ARPA/S30M2/94/01

MRC/AST/FSR/21

ADVANCED SUBMARINE TECHNOLOGY - PROJECT M
PHASE 2: CONTROL EXPERIMENTS AND SIMULATIONS

F.A. Johnson   (Programme Director)

GEC-Marconi Research Centre

1994
Interim Technical Report

19941207 004

ARPA/S30M2/94/01
MRC/AST/FSR/21

ADVANCED SUBMARINE TECHNOLOGY - PROJECT M
PHASE 2:   CONTROL EXPERIMENTS AND SIMULATIONS

Prepared by

F.A. Johnson   (Programme Director)

GEC-Marconi Research Centre
GEC-Marconi Limited
Great Baddow, Chelmsford
Essex, U.K.,   CM2 8HN

March 1994

Accesion For

NTIS CRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	

By _____
Distribution /

Availability Codes

Dist	Avail and / or Special
A-1	

Abstract

The aim of Project M is to demonstrate that the use of broad band active control, with a distributed array of electromagnetic actuators, can not only support a machinery raft but also provide a very high degree of isolation from the hull. In addition the raft can be adequately restrained, to maintain accurate alignment under heavy sea-way or manoeuvring motions, and longer wavelength structural resonances in the raft can be suppressed. This would allow larger lighter rafts can be used to accommodate machinery items without an increase in the acoustic radiation hazard.

This document brings together a description of the work on the control theory, the simulations, the hardware upgrades and the results of the experimental programme which validated the theoretical predictions.

Contents

Figures

Tables

Preface

The work described in this report covers the second phase of Project M. The first phase dealt with the development of the active control theory, and the associated control algorithms, for a very limited frequency range. It was concerned with demonstrating the system identification techniques and with the possibility of suppressing long wavelength structural resonances in one dimension. Simulations were successfully carried out on an idealised raft and an experimental model raft was made to validate this work. The experimental model raft contained 16 electromagnetic actuators which were controlled in the vertical direction only. This was used successfully to demonstrate the ability to perform stable lifts and to carry out, in situ, the system identification tests to define the modal structure of the raft.

The increase in scope for Phase 2 was to extend the active control theory, and its associated control algorithms, to cover the full frequency range, of importance to sea-going vessels, and to include three dimensional effects. It also involved validating this active broad band control theory on an uprated model raft. The uprating of the raft involved converting the existing 16 single axis electromagnetic actuators to tri-axial actuators. This model raft had just sufficient spatial resolution to permit modal control and hence the validation of structural resonance suppression.

This report brings together a description of the work on the control theory, the simulations, the hardware upgrades. It then presents the results of the experimental programme which validated the theoretical predictions.

This page is intentionally left blank.

1 Summary

1.1 Objectives

The traditional method of isolating large machinery installations in ships and submarines has been to mount the machinery on raft structures supported on rubber anti-vibration mounts. For good vibration isolation the rubber mounts need to be soft - but soft mounts result in large misalignment of machinery during heavy weather.

A ship or submarine's machinery raft should be large in order to simplify the installation and minimize the need for flexible connections. However, large lightweight rafts have low frequency structural resonances which rubber mounts just cannot isolate effectively.

Ideally, a ship or submarine's machinery raft should be large and lightweight, without flexing, accurately positioned during heavy weather or high speed manoeuvring, have no mount resonance and yet simultaneously provide excellent vibration isolation at all frequencies. These ideal characteristics are set out in Table 1.1.1 below.

Primary Design Objectives for Marine Machinery Rafts.	
(1)	The raft must be adequately restrained to maintain accurate alignment under sea-way or manoeuvring motions.
(2)	The raft must provide a very high degree of isolation of the hull, over a very wide frequency range, from vibrations due to machinery, or other items on the raft.
(3)	The raft should be large and lightweight so that the overall installation can become much simpler in order to reduce the overall costs.

Table 1.1.1: Ideal Characteristics for Machinery Rafts

Project M seeks to exploit broad band active control, coupled with well designed and well distributed wideband electromagnetic actuators, in order to achieve every one of these ideal characteristics simultaneously.

1.2 Critical Questions

1.2.1 Control Theory and Simulations

In order to achieve these objectives Phase 2 of Project M set out to address three groups of critical questions. The first group of critical questions were concerned with the formulation of the control strategy. These are listed below.

(1)   How is the control strategy translated into practical algorithms.

(2)   How can the different control regimes be moved in frequency by altering the stiffness of the longitudinal raft members.

(3)   How well can the control strategy be tested, on the 16 actuator raft, by this means.

1.2.2 Hardware and DSP Upgrades

The experimental model raft, developed in Phase 1, contained 16 electromagnetic actuators, which were controlled in the vertical direction only over the frequency band 10 to 40 Hz. The increase in scope for this programme, Phase 2, was to upgrade the 16 electromagnetic actuators to tri-axial operation and to extend the frequency range to 0.01 to 500 Hz - the frequency range, of importance to sea-going vessels. It also involved uprating the signal processing system to deal with the increased computational load.

The second group of critical questions were related to these upgrades.

1.2.3 Experimental Programme

The final group of critical questions were addressed by the experimental programme. These are listed below.

(1)   How can one gain a sufficiently stable lift so that system identification tests can be carried out.

(2)   How accurately can the modal parameters be measured and how is this information to be fed to the control algorithms.

(3)   How well and over what frequency band can the raft be supported by a simple modal isolation strategy.

(4)   How will one make the transition from the low frequency modal description to progressively more continuous and more localised higher frequency disturbances when natural damping begins to reduce the coherence length to dimensions less that those of the raft.

(5)   When is it necessary to change from feedback to adaptive feedforward control and how can one effect a smooth transition from one to the other.

(6)   How well can the basic mount resonances be suppressed.

(7)   How well does the sea-way motion control strategy work.

(8)   Can one suppress the longer wavelength bending resonances on the raft using shorter wavelength force distributions.

(9)   To how high a frequency should one aim to suppress internal structural resonances in this way.

(10)   How will one make a smooth transition from this regime to a modal isolation regime.

(11)   Are there any dynamical or numerical stability problems associated with increasing the size of the raft.

(12)   Does mass loading introduce any major problems for the control system.

1.3 Approach

The approach adopted in Phase 2 was to address the three groups of critical questions in turn. Phase 1 dealt with the development of the active control theory and the associated control algorithms for a very limited frequency range. Phase 2 increased the scope by extending the active control theory and its associated control algorithms to cover the full frequency range, 0.01 to 500 Hz. It also involved uprating the 16 actuator model raft by converting the existing single axis electromagnetic actuators to tri-axial actuators, increasing the frequency range of the sensors and uprating the signal processing system to deal with the increased computational load.

The first phase of this programme was the extension of the control theory and its associated control algorithms and the design work needed for uprating the actuators, sensors and electromagnet drive units to cope with this larger frequency range so that an experimental validation programme could be undertaken.

The second phase was concerned with validating the control theory and its associated control algorithms on the uprated experimental facility. The successful validation of a number of key objectives for the basic control strategy should generate the confidence needed to make the critical decision to move into a third, and even more demanding validation phase.

1.4 Results

The experimental results, from Phase 1, showed that one can gain a sufficiently stable lift to carry out the system identification tests in situ and thus determine the modal parameters. These experimentally determined modal parameters were then used for the modal control of the raft over the frequency band 0 to 40 Hz.

The proposal for Phase 2 called for a 10 dB improvement in all the controlled aspects. The results predicted by the simulations and the results obtained in the experiments are listed below.

(1)   How good is the transition from modal to local control?

> This was the most challenging part of the programme. The experimental results depend critically on the specific details of the raft and the low frequency control system. Correlation measurements between the accelerometers and the AC flux loops indicated that the improvement in the isolation of high frequency disturbances at 100 Hz should be between 15 and 35 dB. This result was not verified until after the completion of Phase 2.

(2)   How well is the mount resonance suppressed?

> Simulation result: 20 dB.
>
> Experimental result: 20 dB.

(3)   How well does the sea-way control work?

> Simulation result: 20 dB.
>
> Experimental result at 0.1 Hz: 30 dB.

(4)   Can one suppress longer wavelength bending resonances with shorter wavelength force distributions?

> Experimental result - yes: 20 dB.

(5)   To how high a frequency should this be attempted?

> Experimental result: 20 Hz.

(6)   Are there any numerical or dynamical stability problems?

> The very few numerical stability problems encountered in the experimental programme were all solved. Dynamical stability problems were only encountered at the end of the programme when an attempt was made to see how far the mount resonance frequency could be decreased. The instability arose when the mount resonance frequency was decreased to 2 Hz due to the weak negative spring effect of the analogue flux feedback loop.

3

(7)   Does mass loading introduce any major control problems?

> Experiments with a mass loaded raft performed well as soon as the correct modal parameters were used in the modal control.

1.5 Accomplishments

We have designed and tested very robust, wideband tri-axial electromagnetic actuators, each fitted with sensors to measure position, acceleration and force in the three orthogonal axes. 16 of these actuators have been made and are used, in two rows of 8 actuators, to support a bare model raft structure. The sensor information is fed into a digital signal processing system (DSP), which we also designed and built, and the algorithms in the DSP are used to define the drive currents to the 5 electromagnets in each actuator to generate the required control forces.

We have conducted a major experimental programme which has shown that, although the actual raft dynamics were considerably more complex than that assumed in the simulations, the measured results have been as good as or, in some cases, actually better than the predicted performance.

1.6 Conclusions

These results demonstrate that the control theory has been satisfactorily validated and that the raft components have proved to be satisfactory over the full frequency range 0 to 500 Hz. In particular, they have shown the major advantages that flow from the use of modal control when the spatial resolution of the electromagnetic actuator array makes this possible. Finally, these results show that it should be possible to simultaneously achieve all the ideal characteristics for a machinery raft listed in Table 1.1.1.

1.7 Recommendations

We recommend that this programme should be continued through Phase 3 where the raft will be increased in size so that it becomes flexible in all three dimensions and where realistic machinery distributions can be accommodated on it.

2 Introduction and Background

The powerful propulsion machinery in ships and submarines generates vibrations. These machinery vibrations are transmitted into the surrounding structure, not only causing undesirable noise but also, in extreme cases, resulting in structural fatigue and failure. The traditional method of isolating large machinery installations in ships and submarines has been to mount the machinery on raft structures supported on rubber anti-vibration mounts.

These rubber anti-vibration mounts provide some vibration isolation but they also result in raft and machinery misalignment during the motions of the hull. This misalignment can cause major problems during heavy weather, or high speed manoeuvres, as several hundred tonnes of machinery sways around inside the machinery compartments. For good vibration isolation the rubber mounts need to be soft - but soft mounts result in large misalignment of machinery during heavy weather. To reduce these large misalignments stiff rubber mounts are needed - but such mounts badly degrade the vibration isolation. Poor vibration isolation not only increases the stresses on the ship's hull, but also greatly increases the acoustic noise radiated into the sea. Thus designers are forced to compromise between soft and stiff rubber mounts. The normal compromise is to design the rubber anti-vibration mounts so that the supported raft and machinery have a natural, lightly damped, mount resonance at about 4 Hz.

The reason for this choice can be seen by considering the static deflection of idealised mass-spring isolation mounts. This can be calculated directly from the mount resonant frequency, regardless of the size of the supported mass, and is given by the expression:

$$d \quad = \quad g_o / (2\pi f)^2$$

where g_o is the acceleration of gravity 9.81 m/s/s. The 4 Hz mount thus has a deflection of 16 mm which, under a 45° roll produces an additional deflection of ±11 mm ($\pm d \cos(\theta)$). If the mount resonant frequency was reduced by two octaves to 1 Hz the 45° roll deflection would increase to ±176 mm while if the mount resonant frequency was increased by two octaves to 16 Hz the 45° roll deflection would decrease to ±0.7 mm.

This example illustrates that for good alignment a high mount resonant frequency is required. However, the natural roll-off in the impedance of such passive mounts produces an attenuation floor of 12 dB/octave for frequencies well above the mount resonant frequency. Thus for frequencies of say >30 Hz the vibration transmission through the mount would be increased by 24 dB if the mount resonant frequency was increased from 4 Hz to 16 Hz. Similarly, the vibration transmission through the mount would be reduced by 24 dB if the mount resonant frequency was reduced to 1 Hz. This illustrates why the compromise choice for mount resonant frequencies is about 4 Hz - the roll deflections are just manageable and the vibration attenuation is the best that can be achieved in the circumstances.

A ship or submarine's machinery raft should be large in order to simplify the installation and minimize the need for flexible connections. However, large lightweight rafts have low frequency structural resonances which rubber mounts cannot isolate effectively. This problem is so severe that such rafts are just not a practical proposition with rubber mounts. What is used instead is a number of small, stiff and massive rafts, whose structural resonances are at much higher frequencies where they are less of a disadvantage as a result of the 12 dB/octave roll-off in the vibration transmission. However, such mounting techniques are heavy, bulky and expensive, since installation of each separate item has to be carried out on-board and a large number of flexible interconnections are required.

Ideally, a ship or submarine's machinery raft should be large and lightweight, without flexing, accurately positioned during heavy weather or high speed manoeuvring, have no mount resonance and yet simultaneously provide excellent vibration isolation at all frequencies.

2.1 Broad Band Active Control

The use of broad band active control, coupled with well designed and well distributed wideband actuators, can provide a hierarchy of progressively increasing capability, as indicated below. It is these capabilities that Project M seeks to exploit in order to achieve every one of these ideal characteristics simultaneously.

2.1.1 Control Hierarchy

The features of broad band control we are seeking to exploit can be set out as a hierarchy of progressively increasing sophistication.

Level 1.   Can reproduce the characteristics of individual spring dampers with frequency bands of constant stiffness or constant damping.

Level 2.   Can introduce controllable frequency dependent characteristics to individual actuators.

Level 3.   Can introduce controllable spatial dependent characteristics to arrays of actuators.

Level 4.   Can introduce discrimination effects via the inputs to the control system.

Level 5.   Can introduce discrimination effects via the outputs from the actuators.

Level 6.   Can combine frequency, spatial and discrimination effects to simultaneously satisfy multiple criteria.

Level 1 allows one to make each actuator simply mimic passive spring damper systems.

Level 2 opens up the possibility, for individual actuators, of integrating out low frequency positional errors, introducing damping in particular frequency bands, introducing local nulls to avoid resonances and introducing a progressive roll off in the stiffness with increasing frequency.

Level 3 opens up the possibility, for arrays of actuators, of selectively stiffening or damping low frequency resonant modes or of selectively filtering out higher frequency modes by simply letting them resonate freely.

Level 4 opens up the possibility of discriminating the source of a disturbance so that one can generate different responses for different sources. It also opens up the possibility of discriminating the time duration of different types of disturbance so that one can have different responses depending on the transient or steady state behaviour.

Level 5 allows one to exploit the differing responses of the receiving structures on each side of the actuators to achieve different objectives. For example one can stiffen a raft resonance by using a force distribution which is simultaneously constrained to minimise the acoustic coupling from the supporting structure into the sea.

Level 6 opens up the possibility of satisfying simultaneously a range of different criteria. It gives one a much greater freedom to design the active control system to achieve overall requirements rather than having to make compromises to suit the very limited and fixed capabilities of rubber anti-vibration mounts. It is this level which opens up the possibility of tailoring the dynamics of marine machinery rafts to meet each one of the ideal design requirements simultaneously.

2.2 Uniqueness of Approach

The history of active vibration cancellation began in 1892 when Sir William Yarrow proposed that the vibration in a ship's hull, due to out-of-balance forces generated by the main engine, could be greatly reduced by superimposing an additional vibration force that was exactly equal but opposite. He demonstrated this by attaching two cams, one at each end of the main engine, driving suitable masses to produced the additional vibration forces. When the magnitude and phase of this additional vibration force was exactly equal but opposite to the out-of-balance force generated by the main engine, the vibration in the ship's hull was greatly reduced.

This technique of cancelling single frequency vibrations has advanced rapidly in the last 20 years as modern electronics began to be exploited along with electrically controlled actuators. It is now possible to simultaneously cancel multiple sources of vibration provided each individual source is at a single frequency. The electronics can ensure that the cancelling forces are accurately synchronised and their magnitude and phase automatically adjusted to achieve the best possible performance. In these applications the controlled actuators are used as additional elements designed to make minor improvements to the overall system - they are not used to make fundamental changes to the way the system itself operates.

However, during the 1980's, it began to be recognized that one should be able to exploit electrically controlled actuators to interact directly with the system in order to make fundamental changes to the dynamical properties of the system itself. This new approach aimed to electronically control the dynamics of the system to produce entirely new effects. This is a fundamentally different approach from the synchronous cancellation techniques described above. In demands that the actuators should be able to interact with the system continuously over a wide frequency range - not just at a few selected frequencies. The types of new control that should become possible are briefly outlined in section 2.1.1 above and Project M seeks to exploit all of these new possibilities.

The unique feature of the control exploited in this Project is the use of an array of electromagnetic actuators, with sufficient spatial resolution, to control selectively the main dynamical modes of a machinery raft structure. The advantages that follow from modal control are:

(8)   rigid body modes can be controlled separately from the structural resonance modes. This allows one to heavily damp the mount resonance modes, using a high rate of roll-off of this damping at higher frequencies, without any risk of injecting energy into other modes which could result in instabilities;

(9)   structural resonance modes, which can be resolved adequately, can either be controlled separately or can be released so that they resonate freely without any net force transfer through the mounts;

(10)   structural resonance modes which cannot be resolved adequately, can simply be released to resonate freely;

(11)   when there is redundancy in the spatial sampling of a structural resonance mode the suppression of this mode can be done in ways that also minimize the radiation coupling between the base and the sea;

(12)   the raft can be made to accurately track the rigid body motions of the base at very low frequencies by injecting appropriate forces directly into the corresponding rigid body modes of the raft.

Another unique feature of this system is that all the lift and horizontal positioning forces on the raft are completely under software control. This makes it possible to determine the modal structure of the raft, with any supported masses, in situ.

A key feature of the control used in this Project is the method for linearizing the electromagnetic actuators. It is essential, for accurate control, that every actuator can produce the correct demand force independently of the gap between the armature and the electromagnet. At low frequencies this was achieved by a high gain analogue flux feedback loop.

However, it is also necessary to extend this independence, of force to gap variations, over the full frequency range in order to eliminate any residual force leakage paths. Another unique feature of this Project is the use of a separate high frequency feed forward control algorithm to achieve this result.

Finally, the scale of the electronic control used in this Project is substantially greater than that used in any previous active control work.

3 Objectives and Technical Issues

The primary objective for Project M was to demonstrate that broad band active control, when used with a distributed array of accurate, wide band actuators, could be used to control a marine machinery raft to achieve, simultaneously, the following primary objectives:

(1)   the raft and its supported machinery should accurately follow the motions of the supporting structure so as to maintain accurate alignment under conditions of heavy sea-way or high speed manoeuvring motions;

(2)   the raft must provide the hull with a very high degree of isolation from all types of vibration due to machinery, or other items on the raft, over a very wide frequency range;

(3)   the raft should be large and lightweight so that the overall machinery installation becomes much simpler, in order to reduce the overall costs.

To achieve these three objectives one must also achieve, simultaneously, the following subsidiary objectives:

(1)   eliminate the conventional raft mount resonances;

(2)   eliminate low frequency structural resonances in the raft and its supported machinery. In the special case of machinery rafts for Naval vessels it is possible to eliminate the low frequency structural resonances, using more complex control force distributions, which still further reduce the radiation hazard normally associated with low frequency, long wavelength resonances;

(3)   greatly improve the vibration isolation at all higher frequencies.

The overall objective for this work is to demonstrate that it is possible to tailor the raft dynamics to meet all these requirements simultaneously.

3.1 Phase 1

The previous programme, Phase 1, dealt with the development of the active control theory and the associated control algorithms for a very limited frequency range. It was concerned with demonstrating the system identification techniques and with the possibility of suppressing low frequency long wavelength structural resonances in one dimension only. Simulations were successfully carried out on an idealised raft, of 16 single-axis actuators without any supported masses, and an experimental model raft was made to validate this work over the frequency range 10 to 40 Hz. The experimental model raft contained 16 electromagnetic actuators, which were controlled in the vertical direction only - see Figure 3.1.1. This was used successfully to demonstrate the ability to perform stable lifts and to carry out, in situ, the system identification tests in order to define the modal structure of the raft.

3.2 Phase 2

The increase in scope for this programme was to extend the active control theory, and its associated control algorithms, to cover the full frequency range, of importance to sea-going vessels, 0 to 500 Hz. It has also involved validating this broad band control on an uprated 16 actuator model raft - see Figure 3.2.1. This model raft has just sufficient spatial resolution to validate the suppression of some of the longest wavelength structural resonances in two dimensions. It was also capable of demonstrating the performance possible over a much larger frequency range. Finally, it was able to demonstrate that simple mass loading of the raft did not degrade the performance of the control system.

Figure 3.1.1: Raft for Phase 1

Figure 3.2.1: Raft for Phase 2

The work was split into two parts. The first part involved:

(1)   the development of the control theory, the algorithm structure and simulations of the system performance that can be achieved with an idealised 16 actuator raft with no added masses and with three different longitudinal beam stiffnesses;

(2)   the design of new, or improved components for the electromagnet actuators, amplifiers, sensors and signal conditioning units to deal with the 12 fold increase in the frequency range;

(3)   the design of improved DSP hardware and software.

The second part involved:

(1)   carrying out an experimental and simulation programme to demonstrate the performance that can be achieved;

(2)   the uprating of the 16 actuator raft and all the associated hardware components needed - new sensors, new signal conditioning units and new amplifiers;

(3)   uprating the signal processing system to a maximum of 3.6 GFlops to deal with the increased computational load;

(4)   calibration and commissioning the installation;

(5)   the experimental work to test the control system.

A more detailed description of the work in this phase is given below.

This report deals with each of these topics in turn.

4 The Control Theory and the Simulations

The control theory and algorithm strategy for the low frequency regime had largely been completed. However, the sea-way control algorithm needed to be integrated into the rest of the low frequency algorithm. In addition a satisfactory method needed to be developed for handling the higher frequency regimes. The complete suite of algorithms was then simulated, using an idealised model of the 16 actuator raft with no added masses, to demonstrate that the control algorithms are, in principle, able to achieve the programme goals of a 10 dB improvement in the performance of the active system over that of a 4 Hz passive system. This simulation also considered the effects of changing stiffness of the raft so that different features of the control strategy can be brought out more fully.

The five specific goals for the control theory and algorithms are set out below:

(1)   accurate low frequency alignment (<1 Hz): the ability to maintain constant position of the machinery raft relative to the hull under sea-way and manoeuvring motions;

(2)   the suppression of the fundamental mount resonances (3 Hz - 6 Hz);

(3)   active stiffening of internal raft resonances with simultaneous low frequency isolation (1 Hz - 50 Hz). The raft must provide a stiff platform for mounted machinery, while at the same time providing a high degree of isolation from the hull. This requires the active mounts to be controlled in a collective fashion, taking into account the spatial distribution associated with the modal deformations and the corresponding wavenumber distributions associated with the control force array;

(4)   the isolation of higher frequency internal raft resonances. This requires the active mounts to be controlled in a collective fashion, taking into account the spatial distribution associated with the modal deformations;

(5)   high frequency local mount isolation (>50 Hz). At the higher frequencies of operation, mount vibration transmission becomes more localized and individual mount paths can be regarded as separate, independent sources of excitation. The associated control strategies seek to minimize the vibration transmission through the individual mounts, irrespective of the behaviour of other elements in the active array.

The raft instrumentation, assumed for these simulations, consists of the following components for each of the three axes of an individual mount:

(1)   gap transducers to define the low frequency alignment;

(2)   accelerometers to define the raft dynamical response;

(3)   flux transducers to both control and linearize individual magnet operations and to provide direct measurements of the transmitted forces.

The requirement is to achieve a consistent improvement in both precision of alignment and low frequency isolation compared to conventional passive mounting. The starting-datum is a 4 Hz-resilient mounted passive system for which the corresponding seaway response and force transmissibility characteristics are well-known.

The objective is to seek an order-of-magnitude performance improvement in respect of both low frequency alignment and vibration isolation, relative to the 4 Hz-passive datum. The actual improvement is the subject of the theoretical research and demonstrator program but the contractual requirement is for a consistent 10 dB improvement in every respect.

4.1 Control Theory

The initial plan for the control theory work is listed below.

(1)   The formulation of the control strategy for high frequency isolation.

(2)   The transition from longer wavelength structural resonances to more locally coherent vibrations.

(3)   Establish the techniques for handling banded matrices.

(4)   The transition from feedback to adaptive feedforward control techniques.

(5)   The control strategy for stiffening structural resonances and simultaneously minimising the long wavelength force pattern in the receiving structure.

(6)   The transition from actively stiffening structural resonances to isolating them.

(7)   The simultaneous suppression of sea-way disturbances from below with the isolation of vibration disturbances from above.

(8)   The suppression of mount resonances.

(9)   The transition from sea-way to vibration isolation control.

In the event the control theory development proceeded somewhat differently from the above plan. Items 5 to 9 all relate to various aspects of the low frequency feedback control algorithm while the first 4 items are associated with the introduction of the high frequency feedforward control algorithm. The work proceeded along these two main paths.

These two main algorithms are illustrated in Figures 4.1.1 and 4.1.2 respectively. At low frequencies, the input to the control system is provided by gap and accelerometer signals, and the control system acts directly to suppress fluctuations in these components using a closed-loop feedback system. On the other hand, at higher frequencies the phase lags involved in feeding-back signals become too large for stable operation and a feed-forward open-loop control system becomes necessary.

Figure 4.1.1: Low Frequency Algorithm

13
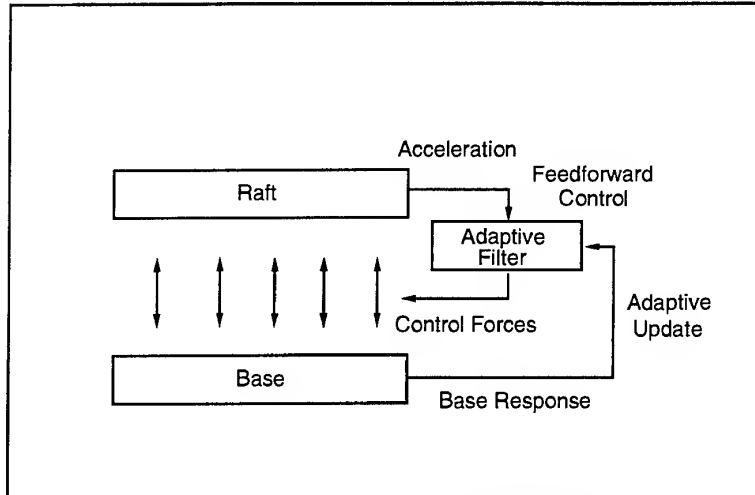
Figure 4.1.2: High Frequency Algorithm

The input to the feed-forward control system is provided by the accelerometers and the output is the control forces, but now these control actions take place at a very low gain, as a consequence of the 12 dB/octave roll-off, so that the control forces have negligible feedback influence on the raft acceleration levels. Thus the objective becomes to define the control system transfer functions to minimise the net response of the receiving hull-structure. This is achieved by performing an adaptive update of the transfer functions as a result of observation of the receiving structure response.

The overall control system is then constructed by implementing these two control strategies in parallel, to give the final structure sketched in Figure 4.1.3. This essentially consists of an inner high-gain loop representing the low-frequency feedback algorithm, and an outer low-gain path which represents the high frequency feedforward process. An important requirement was to ensure that these two processes can successfully coexist, without the action of one compromising the action of the other.

4.2 The Low Frequency Control Algorithm

The low frequency control algorithm detects the raft gap and acceleration signals and outputs the necessary control forces to suppress unwanted low frequency disturbances. As the objective is to isolate the raft structure from the base, these control forces will reduce rapidly in amplitude with increasing frequency, thus giving the controller a markedly low-pass characteristic.

It is also necessary to consider the spatial distribution associated with these control forces so that in the region where they exert a significant restraint on the raft motions, they are arranged to provide the lowest possible coupling interaction with the hull/water radiation characteristics.

This algorithm was specified using optimal state-space feedback control techniques. A feature of this approach is that it provides a "model reference" controller - Figure 4.2.1 - whereby the operation of the physical system is modelled in parallel by the real-time state-space model. At any instant, this model possesses a comprehensive set of estimates of the instantaneous state-variables of the overall system. The control output is derived directly from these estimated states, while the

tracking of the model is maintained by a Kalman Filter update to the state-equations, based on feeding-back the difference between the predicted outputs of the model and the actual observed system outputs.
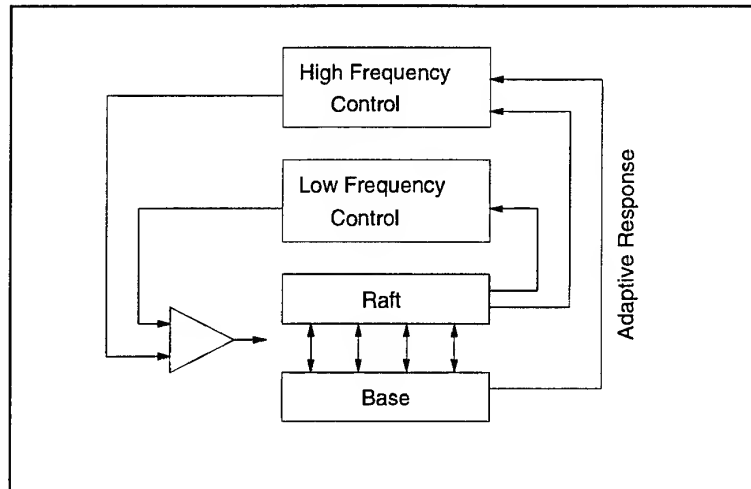
Figure 4.1.3: Overall Algorithm
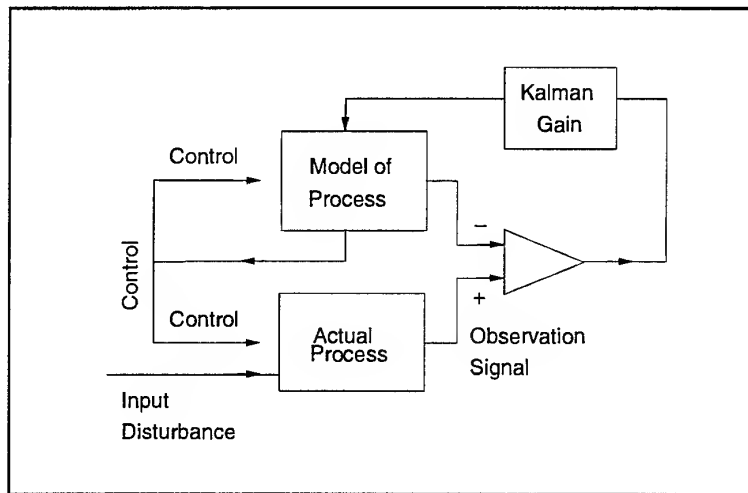
Figure 4.2.1: Model Reference State-Space

There are several advantages associated with this approach. Firstly, once the system model has been accurately established, the technique permits rigourous, stable optimal feedback control to be applied, since the control can be derived directly from each and every (estimated) state of the system. Secondly, the associated feedback of the Kalman filtered observations can be specified according to optimal criteria for the rejection of unwanted noise in the observations. Finally, the

existence of a parallel system model enables a continuous comparison to be made between the behaviour of the model and the behaviour of the actual system, thus providing a basis for future adaptive modification of the model, together with capability for rapid detection of transducer faults in the real system.

The approach used was to decompose the motions of the structure into a finite set of modes. The corresponding states were then defined as the amplitudes and velocities of each mode respectively. The state-space was than applied as a set of fully independent modal control actions. However, the rigourous solution required additional criteria to be satisfied. The specific constraints imposed by these optimality conditions lead to a degree of cross-coupling between the modes. Fortunately, however, it is found that the rigid-body motions of the assembly can always be decoupled from the flexural motions, so that the problem can proceed with separate conditions for the rigid body and flexural components of motion.

The optimality conditions have to summarise the designer's requirements in an appropriate set of quadratic cost functions. These requirements can be most conveniently introduced as an additional space/time filtering action on the actual system observation and control variables, so that the output of this filtering action defines the various space/time weighted performance criteria the designer wishes to optimize. When this filtering action is included in the overall definition of the system model, the quadratic cost functions for the modified output signals take a very much simpler form, and the subsequent solution of the problem becomes relatively straightforward.

This approach is illustrated in Figure 4.2.2. The control and observation parameters are passed into a filtering process which generates a new set of control and observation components. This transformed set of components is then incorporated into the overall system state equations. The control now has to minimise the mean square sum of the filtered observation and control signals and the designer can therefore concentrate his effort on specifying this space/time filtering process to satisfy his particular objectives, without undue concern for the subsequent solution procedure.

A brief summary will now be given of the typical space/time filtering processes which are required, with reference to Figure 4.2.3. The objective is that the outputs of the filtering process should represent the significant contributions to the cost function.

4.2.1 Gaps Control

The objective is to ensure that at very low frequencies the actual gaps are accurately aligned, but this requirement is relaxed as the frequency increases and the raft is then allowed to vibrate freely. Thus the corresponding modal gap outputs are low-pass filtered with a cut-off around 1Hz, so that these output components assume significant values only at very low frequencies.
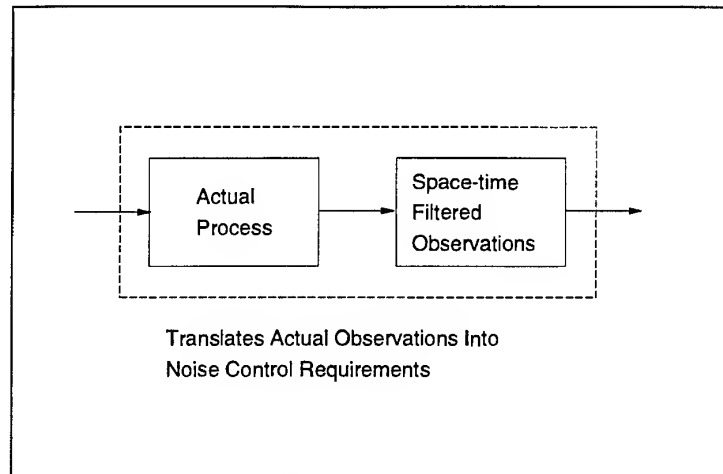
Figure 4.2.2: Modified Observation Process
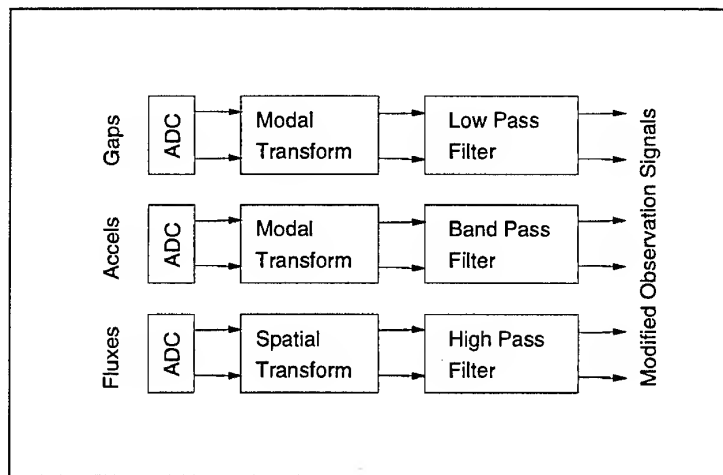
Figure 4.2.3: Detail of Observation Process

4.2.2 Non-Reciprocal Criterion

The requirement is to ensure that those rigid-body accelerations which are necessary to follow hull seaway motions are accurately driven by the control forces, while rigid-body accelerations which are induced by vibration excitation are ignored, and excite no corresponding force action. This requirement can be met by imposing a combined constraint on the rigid-body modal gap and acceleration components, namely

$$\text{Minimise } \{ F - M (a - \ddot{g}) \}^2$$

where the F is the control force, g is the rigid-body modal measurement gap, and a is the rigid-body raft modal acceleration component.

Under circumstances where the motions are entirely due to vibration, the gap derivatives g and raft acceleration components a are equal; the cost function then seeks to drive F to zero. However, when the motions are due to seaway demands, the function (a − g̈) effectively measures the hull acceleration. The overall cost function is then at minimum when the control force is of precisely the size necessary to drive the raft with an acceleration matching that of the hull.

4.2.3 Flexural Mode Control

The requirement is to ensure that low frequency internal resonances of the raft structure are suppressed, so the corresponding acceleration components are band-pass filtered. Each filter is centred on the resonant frequency of its respective mode, so that its contribution to the cost function is greatest at its resonant frequency. For higher order modes, modal accelerations are allowed to run free, and the corresponding band-pass filter gains are set to zero. So the components of acceleration associated with the higher order modes play no part in the control response.

4.2.4 Spatial Force Distribution

An important requirement is that the control actions which are brought into play to suppress the low frequency resonant modes should possess spatial distributions which minimise the efficiency of transmission into the supporting hull structure. A regularly distributed control force array can be decomposed into spatial spectral components, simply by imposing a discrete spatial Fourier Transform operation. The corresponding transform components are referred to as spatial wavenumber components, and it is well-known that it is the long wavelength, low-wavenumber components which transmit most efficiently into the hull receiving structure and the sea. In contrast, the troughs and crests of the shorter wavelength high wavenumber components tend to mutually interfere, giving rise to a spatial cancellation effect. As a result, they transmit much less efficiently into the receiving structure.

Thus when seeking to control the internal resonances of the raft structure, it is preferable to do so wherever possible with short wavelength, rather than long wavelength control force distributions. This in turn can be quantified by imposing suitable constraints on the choice of wavenumber component; the low wavenumber force components carry a high penalty, while the high wavenumber force components carry a much lower penalty.

The control force array can therefore be transformed using an orthogonal spatial transform, to define an equivalent set of controls in the spatial transform domain. The cost penalty is then applied most severely to the low wavenumber components. This constraint takes the form of a set of conceptual high-pass filters cutting on at progressively higher frequencies. Thus, for the lowest wavenumber component, the high pass filter comes into play at a very low frequency, whereas for the higher wavenumber components the cut-on becomes effective at correspondingly higher frequencies. In general, these cut-on parameters should be chosen to mirror directly the transmission and radiation cut-on characteristics of the corresponding wavenumber patterns.

The above discussions relate to the cost penalty which should be applied to the observation of the control action, namely the control fluxes. However, in practice it is equally possible to impose the equivalent cost constraint directly on the demand for control action.

But in the latter context, it becomes easier to choose the central mathematical set of control parameters as those variables which would correspond to the outputs of this high-pass filtering process; the real control signals can then be derived by low-pass filtering and inverse spatial transforming these central control parameters. The necessary low-pass filters become naturally defined as the complementary inverse of the corresponding high-pass filters. This feature is illustrated by transforming Figure 4.2.3 to the equivalent representation of Figure 4.2.4.1. It can be seen that the model of the complete system now incorporates both front-end and back-end filtering and transformation processes.
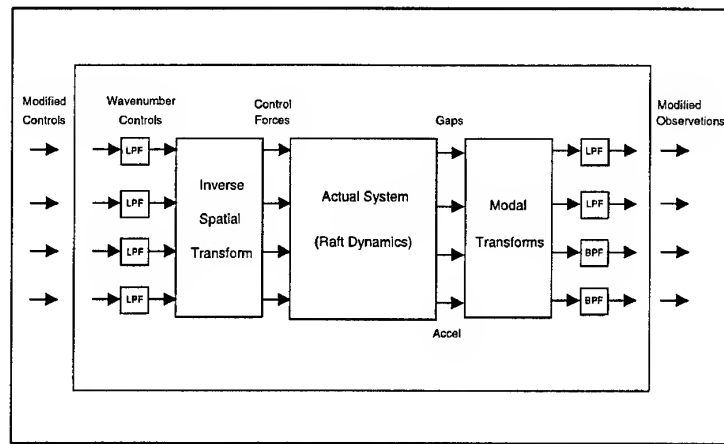
Figure 4.2.4.1: Front and Rear End Processes

The advantage of this latter approach is that all the filters which are actually implemented become low-pass components, and so are much less sensitive to high-frequency noise. The net cost-function can then be applied as an equal weighting against the central control components; the subsequent low-pass filtering to derive the final control outputs will automatically suppress the low-wavenumbers and bias the outputs in favour of the high-wavenumber components.

4.2.5 Decoupling of Rigid Body Motion and Flexural Modes

It has previously been remarked that for simplicity one would wish to control each mode of motion of the structure independently, but the imposition of the cost constraints leads to cross-coupling between the different modes of motion. This cross-coupling arises as a consequence of the dual orthogonal structure of the control/optimization problem. The modal decomposition of the structure defines one complete set of orthogonal eigenvectors which can describe fully any force distribution. But the spatial decomposition of the control force array provides a second complete set of orthogonal eigenvectors which can also fully describe any force distribution. These two orthogonal sets are not the same - the cross-products between the two sets do not vanish.

It is this feature which enables alternative choices of control action to be made. For example if one wishes to control a specific resonance mode has a choice of using a force distribution with either a low-wavenumber component or with a higher wavenumber component.

The purpose of the optimization is to determine the specific choice of control action which can adequately leverage a given dynamical mode while minimising the spatial cost function; the latter function is defined in terms of the spatial eigenvectors.

Although these two orthogonal sets of vectors are not the same, they can be arranged to share one set of eigenvectors in common, namely the rigid body modes. These components of motion are always orthogonal to the higher bending modes of the structure, while momentum considerations ensure that the rigid body modes will not be excited by any force distribution which has zero net linear or angular momentum. So neither a higher flexural mode of the structure, nor a zero-momentum wavenumber mode of the spatial force distribution can influence the rigid body modes.

As a result, the optimality conditions for the rigid body motions of the uniform raft can be fully decoupled from both the flexural structural motions and the zero-momentum wavenumber components. Thus, the control problem can be separated-out into two distinct problems, namely the optimal control of the rigid body modes, and the optimal control of the flexing modes.

While this separation of the rigid body modes from the flexural modes does not in any way modify the general philosophy of the control problem as previously set out and described, it does make for a useful simplification in the present numerical calculations. Instead of handling a single large matrix optimization, the practical optimization can be performed as two separate, non-interacting smaller calculations. This has been found to substantially reduce the sheer size of the numerical arrays which are required for each specific calculation, and also to improve the speed of convergence of the associated numerical solutions.

4.3 High Frequency Control Algorithm

The high frequency control algorithm operates according to an entirely different philosophy from the low frequency algorithm. Whereas the low frequency algorithm makes use of substantial feedback gain to suppress unwanted disturbances, the high frequency algorithm operates in a feedforward mode with gains which are generally of order unity or less. Moreover, whereas the low frequency algorithm operates in a global sense forming each control output from a linear combination of all the sensor inputs, the high frequency algorithm operates entirely local to each individual mount, sensing the immediate input acceleration to the mount, and driving the control electromagnets to minimise the force transmission into the base structure.

Since this control process operates at low gain in an essentially open loop fashion, it is necessary to use a secondary observation process to monitor the overall performance and adjust the control parameters to achieve the desired response. Further it is necessary to ensure that operation of the high frequency algorithm does not interfere with and compromise the low frequency algorithm, which it is assumed has already been optimized over its working frequency range.

A feature of such adaptive feedforward processes is that the processing workload is generally more intensive than the corresponding demand for a feedback controller. This is because all the details of the transmission process must be reproduced accurately within the feedforward path; the requirement for precision in this respect is much tighter than for a conventional feedback process. A typical requirement is to reproduce features of detail with 2Hz resolution using a 2kHz sampling rate (500Hz bandwidth). This necessitates a 1000-point adaptive FIR filter, with a corresponding workload of 4MFlops for the baseline filtering requirement, and 8Mflops if the adaptive update is to be maintained every sample.

One can summarise the following requirements which the algorithm must satisfy.

(1)   It must seek to reproduce unwanted transmission into the base structure (eg arising from eddy current and local air compression effects) and regenerate this in anti-phase, thus cancelling these components of transmission.

(2)   It must possess an appropriate method of performance monitoring and optimization, to maintain the net transmission process at the least possible level.

(3)   Although individual mounts are controlled on a local basis, the subsequent base structural response is due to the combined effects of several transmission paths. An appropriate deconvolution process must be implemented to ensure that the minimisation of each transmission path takes place without cross-interaction.

(4)   The optimization process must take place over the higher frequency regime, without compromising the existing optimized performance associated with the low frequency algorithm.

The high frequency algorithm address all of the issues (1)-(4) above. It must be noted, however, that in certain respects the interaction with the low frequency algorithm is already minimal, since the latter is a global algorithm while the high frequency algorithm is entirely local in its operation.

The advantage of this feature is that the interaction between the global low frequency and local high frequency algorithms will be correspondingly reduced, but the corresponding disadvantage is that the high frequency algorithm will be unable to "make up" for any poor design in the roll-off associated with the low frequency algorithm. Thus the rate of roll-off associated with the global algorithm defines a lower bound on the overall performance, and the purpose of the high frequency algorithm becomes to ensure that this lower bound is indeed achieved, and is not compromised by local leakage paths associated with the mount transmission.

The high frequency algorithm was implemented, in the first instance, by using an appropriate modification of the standard "filtered-x" algorithm. The central feature of this algorithm is that the transmission process from the controller output to the observation points is simulated in parallel (the filtered-x component). Any residual error at the downstream observation locations can then be correlated with the response which would have been expected at these locations, given the known signals which have been applied to the actuators.

Any residual correlation is then used as a correction to the actuator drive filters, in order to drive the correlation to a minimum. An important modification to this algorithm in the present context is that it includes the cross-correlation of the responses at all the downstream observation points, thus enabling the cross-interactions between the different actuators to be taken into account.

As already remarked, it is essential that the algorithm should only modify the high frequency components of the actuator drive signals, and should not seek to modify the existing optimized low-frequency behaviour. In this respect, it is not sufficient simply to high-pass filter the information supplied to the algorithm, since this leaves the algorithm free to assume arbitrary values at low frequency. A positive constraint must be applied to the low-frequency behaviour.

This constraint is achieved by simulating the known effects of the low frequency algorithm, and subtracting these off from the error correlation process. This ensures that the low-frequency operation is left intact, and that non-zero error correlations relate only to the high frequency behaviour. More importantly, since the information supplied at the input to the algorithm is now all-pass, any attempt to modify the low-frequency behaviour will definitely register as an "error", so the constraint is imposed in a positive fashion, rather than simply by default.

4.4 Basic Feedforward Adaptive Algorithm

The basic operation associated with an LMS adaptive algorithm is shown in Figure 4.4.1 The objective is to control an FIR filter so as to reproduce in anti-phase a given process, by modifying the filter tap weights (coefficients) in response to the observed performance error. The simplest form of adaptive update algorithm is the Widrow algorithm, which updates the array of tap-weights based on the instantaneous cross-correlation between the filter input data array $X(t)$ and the observed error $\varepsilon(t)$.
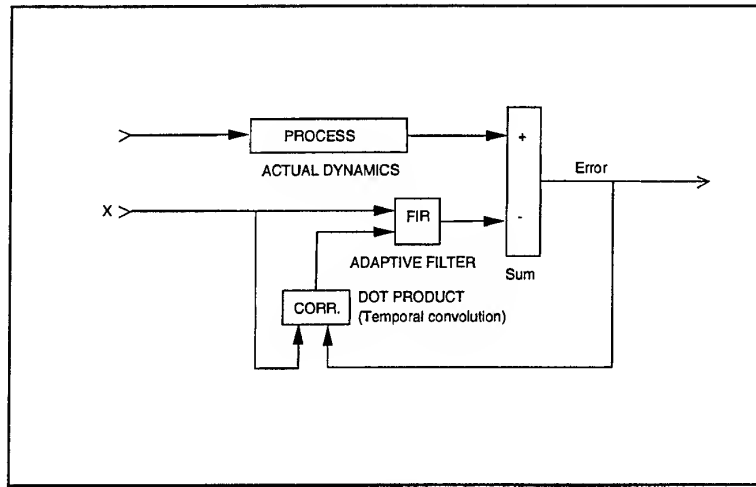
Figure 4.4.1: Basic Feedforward Algorithm

The objective is to minimize $\overline{(y - F \cdot X)}^2$ by the choice of filter coefficients F noting that:

$$\nabla_F (y - F \cdot X)^2 \quad = \quad -2(y - F \cdot X)X$$

and that the right hand side of this equation is the instantaneous gradient which can be interpreted as the cross-correlation between the filter input data X and the observed error ε where:

$$\varepsilon \quad = \quad (y - F \cdot X)$$

The optimal performance is achieved when all components of the cross- correlation are driven to zero, so the adaptive update seeks to modify each tap by an amount proportionate to any residual non-zero correlation, that is:

$$\delta F \quad = \quad 2\mu\varepsilon X$$

4.5 Downstream Propagation and Interaction

In practice, observation cannot take place immediately downstream of the adaptive filter, but rather after the combined signals have propagated through a further downstream process. Thus the x-input to the filter is no longer accurately time-aligned with the observed error, and steps must be taken to realign this data. This can be achieved using the "filtered- x" algorithm, which passes the x-input data through a simulation of the downstream process, prior to performing the correlation update. This restores the alignment, and ensures that the residual error is accurately correlated with the expected response at the observation point see Figure 4.5.1.
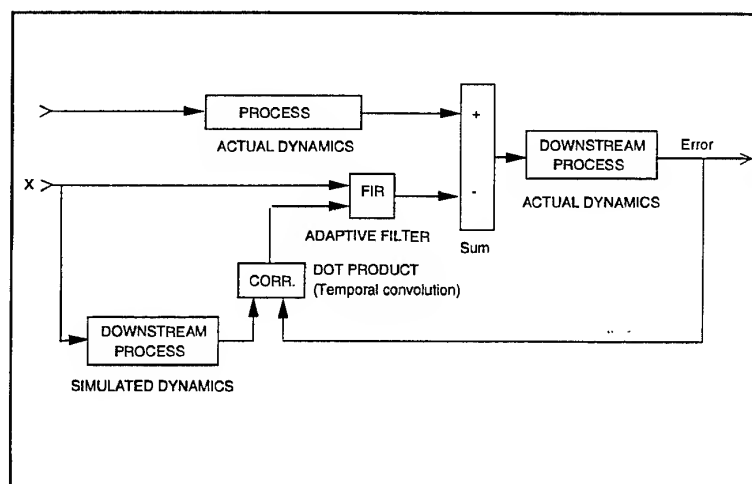
Figure 4.5.1: Filtered-X Adaptive Algorithm

This additional simulation requirement may itself be implemented using a separate adaptive filter; this involves injecting test signals into the downstream process, and observing the associated response at the observation point. An adaptive filter can be used to characterise this transmission directly. Once it has adequately converged, it can be "frozen" and inserted as the downstream simulation element of Figure 4.5.1.

4.5.1 Deconvolution of Downstream Interaction Paths

The algorithm described above enables the effects of temporal misalignment arising from downstream propagation to be compensated. But in the present context one must also take account of the fact that the error responses are induced by a multiplicity of separate mount-transmission paths. Thus in practice, one seeks to update a parallel array of adaptive filters, each simulating a separate mount-transmission path, by observing a corresponding array of error responses.

Under such circumstances, one must implement the correlation process with respect to the spatial characteristics of the observed response, as well as its temporal behaviour. Effectively, in order to minimise the response due to any specific mount path, one seeks out those components of response which simultaneously exhibit the correctly matching spatial and temporal pattern.

This can be achieved using the multi-degree-of-freedom filtered-x algorithm as shown in Figure 4.5.1.1. Each input to each adaptive filter is passed through a simulated downstream process, in order to generate an array of outputs corresponding to the responses which would be expected at the various observation points.
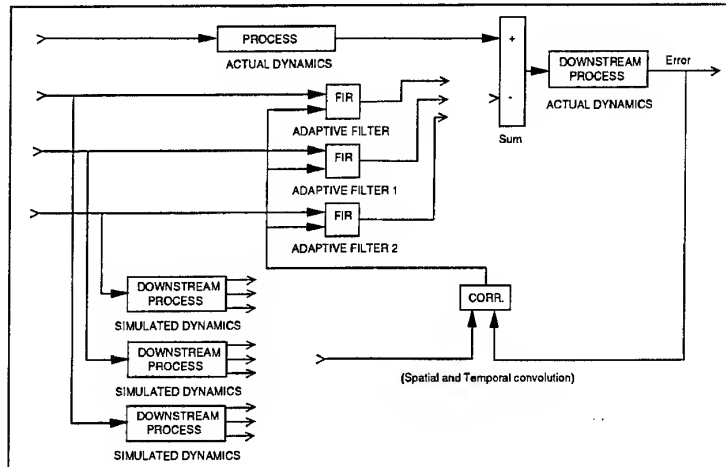
Figure 4.5.1.1: Multi-Degree-of-Freedom Adaptive Algorithm

Each individual array of simulated responses is then correlated with the full array of observation error-responses. This defines a set of scalar products of space-time correlations, which are then used to define the coefficient updates for each corresponding adaptive filter.

It can be seen that the computational workload associated with this adaptive update can very rapidly exceed the workload within the primary filtering process, since the simulation of the downstream array dynamics involves a multiplicity of cross-paths. For example, for a 48-input, 48- output system, there are 48 feedforward paths, but depending on the severity of the cross-interactions, the downstream filtered-x process can involve simulating up to 48 x 48 transfer functions.

It should be noted, however, that various simplifications can be applied to this downstream process. In particular, if the updates are not applied on a sample-by-sample basis, but rather in sequential rotation, the workload can be substantially reduced (at the expense of longer convergence times). Moreover, many of the downstream processes share common dynamical features (eg resonant frequencies), so it can be more economic to implement these processes using IIR (infinite-impulse-response) filters, with common poles.

But the most effective technique is undoubtedly to try to ensure that each downstream observation only responds to a local subset of inputs, thus enabling the 48 x 48 matrix array of transfer functions to be condensed to a banded or near-diagonal array.

In the present context, it is proposed to achieve this by using the magnet flux response as a measure of vibration transmission. Clearly, the flux induced within any particular mount is not sensitive to the action of other mounts, whereas the vibration response at the base of each mount will be influenced by the combined actions of all the mounts.

For this reason, each mount is instrumented with passive flux-sensing coils, intended to provide clean, low-interaction measures of the residual transmitted forces.

Thus, it is intended that the downstream process should be simplified to yield a diagonal or near-diagonal form with only local interactions. However, the practical extent of such interactions can only be determined by experiment. Nevertheless, the algorithm structure presented in Figure 4.5.1.1 can in principle handle any such level of interaction.

4.6 Interaction with Low-Frequency Algorithm

Since the low frequency algorithm has been carefully designed to achieve optimal performance while satisfying a number of constraints, one does not wish the high frequency algorithm to introduce subsequent low frequency perturbations which compromise this careful optimization.

For this reason it is not sufficient simply to high-pass filter the data which is supplied to the adaptive algorithm, since this will mean that there is no information at all to constrain its low-frequency characteristics; the latter will then be established at random according to the specific circumstances. Indeed the algorithm will almost certainly seek to equalise the effects of any high-pass filters, and will seek to boost its low-frequency gain in order to do so. So it is important that relevant low-frequency information is still made available to the algorithm, in such a manner that if it departs from its constrained characteristic, errors are registered and positive corrections are applied.

This requirement can be achieved using the circuit of Figure 4.6.1 It can be seen that this is a basic filtered-x algorithm, but the output of the adaptive filter is split into complementary high-pass and low-pass components (the sum of these complementary filters is unity). While the high-pass component is then injected into the actual system, the low-pass component is passed through a parallel simulation of the downstream process, and then subsequently subtracted-off from the error-sensor.

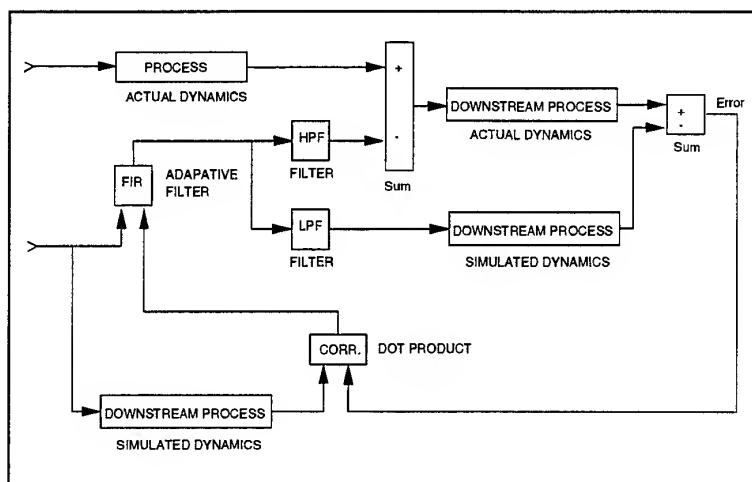**Figure 4.6.1: Adaptive Algorithm with Low Frequency Constraint**

Providing the simulation dynamics is an accurate representation of the actual downstream process, the overall dynamics of this circuit are identical to those of a standard filtered-x adaptive algorithm. The fact that the output is split into two complementary high-pass and low-pass components does not affect any of the data supplied to the algorithm.

The adaptive filter will then correctly seek to reproduce both the low-frequency and high-frequency characteristics of the desired process. But only the high frequency component, with controlled high-pass cut-on, is introduced into the real cancellation process. The low-pass components are kept out of the real system, but they are re-introduced into the error path prior to the update.

So the error-update is an all-pass process, which seeks to maintain accurate alignment over the full frequency range, but only the constrained high frequency components are introduced into the actual dynamics of the real system.

Having thus identified an algorithm which will achieve the desired objective, it is appropriate to try to simplify it.

In Figure 4.6.2 this is achieved by noting that the parallel low-pass components could equivalently be derived from the true signals emerging from the actual process, rather than from the output of an adaptive filter which is seeking to reproduce this process. This relieves the adaptive filter of the need to converge so accurately at low frequencies, since the requisite low-frequency component of output data is now automatically available to full precision.
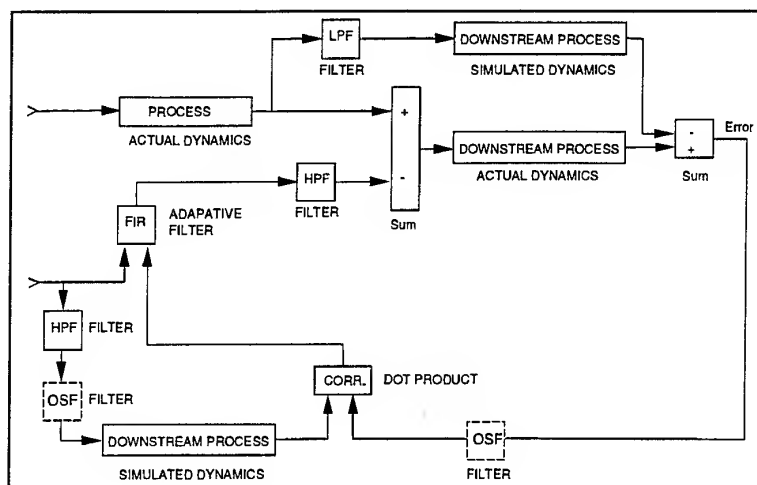
Figure 4.6.2: Adaptive Algorithm with Simplified Low Frequency Constraint

The requirement now becomes that the adaptive filter should accurately converge at high frequencies; moreover it will no longer seek to compensate for, or "equalise", the effects of its own output high-pass filter, since the overall process which it must now reproduce is also a matching high-pass process. Yet, if any significant low frequency error is generated, this will propagate through to the error observation point, where it will be taken into account and corrected in the correlation update process.

In its present form, this algorithm requires a matching high-pass filter to be placed in the filtered-x portion of the circuit. But this filter roll-off could be flattened-out or removed completely, provided a compensating low-frequency weighting correction is applied to the error-output path as indicated. The use of such a filter would recover the all-pass characteristics of the adaptive update, and would serve to tighten-up still further the low-frequency precision of operation.

4.7 The Overall High Frequency Algorithm

The overall algorithm for the system is set out in Figure 4.7.1.

It can be seen that the central portion of the circuit contains the primary low-frequency feedback algorithm, which has been described earlier. For the high frequency algorithm, the acceleration response of the raft is detected, and used as input to the adaptive FIR filters which simulate the mount transmission characteristics. The outputs of the FIR elements are high-pass filtered and fed-forward for addition into the active mount drive signals.

The adaptive update is performed by applying a matching high-pass filter to the acceleration response, and feeding it into a simulation of the dynamics of the base supporting structure. This simulation calculates the separate response at each of the observation points which would be generated by the individual accelerometer signals, and correlates this information with the actual error signals observed in practice. This correlation defines the adaptive update for the FIR filters.

The low-frequency constraint is applied by passing the actual low-frequency control signals through a separate simulation of the base dynamics, and subtracting the simulated response signals from the actual observed response of the base.

Figure 4.7.1: Overall High Frequency Algorithm

In this form the algorithm contains all the features which have previously been identified separately. Specifically, the downstream response is simulated by a set of filtered-x paths, the correlation is performed taking into account spatial cross-interactions in the response, and the algorithm is constrained so that it does not seek to modify the existing low-frequency operation of the system.

It should be noted that specification of the separate components in the algorithm is largely defined by the dynamical characteristics of the actual installation, so there is little scope for ambiguity in this respect. The principal variable which can be chosen independently is the exact specification of the complementary high-pass/low-pass filters. The characteristics of these filters will define the precise cut-on point of the algorithm, while the precision with which these filters approximate to unity in their pass-band will define the limiting precision of the algorithm.

4.8 Simulations

The control algorithm structure was then tested in a series of simulations to address the following critical questions.

Critical Questions.

(1)   How is the control strategy translated into practical algorithms.

(2)   How can the different control regimes be moved in frequency by altering the stiffness of the longitudinal raft members.

(3)   How well can the control strategy be tested, on the 16 actuator raft, by this means.

The simulations were carried out to demonstrate how these control algorithms would work with an idealised 16 actuator raft. The detailed tasks carried out are listed below.

(1)   Carry out simulations on an idealised raft to demonstrate the following:

   (i)   sea-way control;

   (ii)   mount resonance suppression;

   (iii)   structural resonance suppression - modal;

   (iv)   structural resonance isolation - modal;

   (v)   high frequency isolation of localised vibrations.

(2)   Repeat these simulations for several different raft stiffness configurations.

The performance of the broad band active control strategy was directly demonstrated by making a comparison with the performance of a purely passively mounted raft with a 4 Hz mount resonance frequency. The expectation was that one would achieve performance improvements of the active system relative to that of the passive system of at least 10 dB. The simulations also demonstrated important features such as the mean acceleration on the raft and the mean vibration transmission, both of which were compared with the corresponding figures for the passive case.

4.9 Results and Discussion

Three different raft models were simulated: one, the "soft" model with longitudinal resonant frequencies in the range 3 - 20 Hz; one, the "medium" model with longitudinal resonant frequencies in the range 6 - 40 Hz; and one, the "stiff" model with longitudinal resonant frequencies in the range 15 - 200 Hz. The detailed results and discussion are listed in the document MRC/AST/FAR/11.0. The following gives a brief summary of the main findings.

4.9.1 Raft Responses

The first set of results show the performance for the three raft models when the control algorithms were simulated, in detail, as a series of time-stepped models and the performance data analysed to yield the frequency distribution of the mean square acceleration, in response to a vibration excitation from a random point source, for the three models. These results are shown in Figures 4.9.1.1, 4.9.1.2 and 4.9.1.3. Each figure shows the results for an optimal low-frequency control compared with the response for passive mounts.

Figures 4.9.1.4, 4.9.1.5, and 4.9.1.6 show the same frequency distributions of the mean square acceleration, in response to a vibration excitation from a random point source, for the three models when the calculations were performed directly off-line and in consequence these results are more accurate than the results in Figures 4.9.1.1, 4.9.1.2, and 4.9.1.3 where the accuracy was limited by the computation time used in these simulations. Nevertheless the comparison between these two sets of results shows that good agreement is obtained.

Figure 4.9.1.1: Soft Raft Response (Real Time)

**Figure 4.9.1.2: Medium Raft Response (Real Time)**

Figure 4.9.1.3: Stiff Raft Response (Real Time)


Figure 4.9.1.4: Soft Raft Response (Off Line)

**Figure 4.9.1.5: Medium Raft Response (Off Line)**

Figure 4.9.1.6: Stiff Raft Response (Off Line)

It is immediately apparent that in the case of the actively controlled soft raft, Figures 4.9.1.1 and 4.9.1.4, the first 3 low frequency raft resonances are completely suppressed, and the mean acceleration response approximates nearly to a straight line. This means that the raft is behaving almost as a rigid body with virtually no evidence of low-frequency resonances. In contrast, it can be seen that the response of the passively mounted structure is extremely resonant, with amplifications of up to 20 dB over the actively controlled structure. This corresponds to motions which are a factor of 10 greater in amplitude. The 4th resonance is not suppressed, although it is still at a comparatively low frequency (20 Hz), but this is primarily due to the fact that there are only 8 mounts supporting the structure.

Similar results are obtained for the medium raft, Figures 4.9.1.2 and 4.9.1.5, but in the case of the stiff raft, Figures 4.9.1.3 and 4.9.1.6, the control system can only partially suppress the first resonance, and thereafter the response is very similar to the passive raft. The reason for this is that the suppression of resonances in the stiff raft requires much larger forces to be exerted at higher frequencies, and the active mounts cannot supply such forces without compromising the isolation of the structure from the hull. Thus it is better to "let these resonances go". On the other hand, for the soft raft at very low frequencies, the active mounts can indeed supply the requisite control forces while still maintaining overall isolation of the structure.

It is clear from these results that for resonant frequencies up to 30-40 Hz, very effective suppression is possible, but at higher frequencies it is better not to suppress the structural motions. The active control strategy is most effective on flexible structures, which have a number of low frequency resonances, but is less effective on stiffer structures. However, this is the feature that one wishes to exploit in order to make large, lightweight extended machinery raft installations practical.

4.9.2 Radiated Noise

In order to quantify the potential improvements radiated noise to define a typical cost function associated with each wavenumber distribution. The cost function used is shown in Figure 4.9.2.1. The lowest wavenumber is assumed to couple efficiently (ie 100%) over the full frequency range 0-100 Hz, but the higher wavenumbers are assumed to have linear cut-ons, each starting at a progressively higher frequency as indicated. This is typical of the radiation characteristics of the higher-order spatial wavenumber distributions.

The results of these radiation efficiency calculations are shown in Figures 4.9.2.2 and 4.9.2.3 for the soft and medium raft models. It is apparent in both cases that the transmission efficiency is substantially reduced compared to the equivalent passively mounted structure.

4.9.3 Seaway Results

As the rigid body modes of all three raft models are the same their response to sea-way disturbances will also be the same. A typical seaway disturbance is shown in Figure 4.9.3.1. This shows a hull input motion representing a peak displacement of 5 metres, and an average displacement of 1-2 metres.

The response of a passively mounted raft relative to these hull motions would amount to some 15 mm peak, with average displacements of 3-5 mm. Figure 4.9.3.2 shows the actively mounted raft shows displacements which at all times are within 1 mm, and generally considerably less. Thus it is clear that the actively controlled raft can track and follow the hull motions extremely accurately.

Wavenumber Cost Function

Figure 4.9.2.1: Wavenumber Cost Function

Radiated Noise: Soft Raft

Figure 4.9.2.2: Soft Raft - Radiated Noise

Radiated Noise: Medium Raft

Figure 4.9.2.3: Medium Raft - Radiated Noise

Figure 4.9.3.1: Seaway Input Disturbance

Figure 4.9.3.2: Raft Seaway Response

Figure 4.9.4.1: Soft Raft - High Frequency Performance

4.9.4 High Frequency Results

These simulations assumed that the performance would be monitored by measuring residual force transmission, using the flux loops, and that these sensors would not be subject to cross-interactions. The assumption was also made that each mount was subject to a component of unwanted leakage flux variation due to the local vibration levels associated with each magnet armature. This situation is very likely to arise in practice.

The high frequency algorithm was then run in parallel with the low-frequency algorithm. The excitation inputs were taken to be a distribution of 16 random, uncorrelated forces, so that the corresponding motions of the raft structure were of a completely random nature, with poor spatial correlation at the higher frequencies. The criterion which was adopted for assessing performance was the summed-square force transmission through all the mounts. Figures 4.9.4.1, 4.9.4.2 and 4.9.4.3 show the calculated summed-square transmitted force for each of the three raft models, before and after adaptive cancellation.

Figure 4.9.4.2: Medium Raft - High Frequency Performance

It will be seen that in each case the adaptive control process can cancel out the unwanted local components of flux transmission. The overall transmitted forces are substantially reduced to a lower bound which is defined by the residual forces associated with the low-frequency algorithm. An interesting feature is also observed in Figure 4.9.4.3, namely that the response for the stiff raft contains resonant peaks. This shows that local mount leakage transmission can excite structural resonances. The adaptive algorithm can knock out these resonant peaks and reduce the net transmission to that defined by the low-frequency, global controller.

36

Figure 4.9.4.3: Stiff Raft - High Frequency Performance

4.10 Conclusions

These simulations assumed that the raft was pin jointed and that the forces would be applied directly at the pin positions - a rather idealised model. These simulations indicated that the methods proposed for the broad band active control of a large, lightweight flexible raft should be able to achieve the following results for the bare raft:

(1)   20 dB suppression of the response to sea-way motions;

(2)   the elimination of the mount resonance;

(3)   20 dB suppression of lowest frequency raft resonance;

(4)   40 dB suppression of the acoustic radiation associated with the lowest frequency raft resonance by the use of more complex control force distributions;

(5)   10 dB to 40 dB attenuation of higher frequency (above 40 Hz) vibration transmission.

In every case the simulation indicates that the target improvement of 10 dB can be substantially exceeded. It now remains for the experimental programme to see to what extent these simulated results can be realised in practice.

5 The Raft Hardware

The raft that was developed during Phase 1 had vertical lift actuators only which were powered by pulse width modulation (PWM) amplifiers. Each actuator had 3 sensors: a low frequency Hall effect flux sensor, a proximitor for measuring the magnet gap and an accelerometer to measure the vertical acceleration of the raft at the actuator. As the frequency band of interest was 10 to 40 Hz and the objective was to demonstrate the control of rigid body and low frequency flexural modes only, the amplifiers were not optimized for low noise and the sensors had low bandwidths.

The Phase 2 programme required that the actuators be upgraded to tri-axial actuators by the addition of 4 lateral electromagnets - operating as opposing pairs - to each actuator, the amplifiers powering the vertical lift electromagnets to have their noise levels reduced by at least 30 dB, suitable amplifiers developed for the lateral electromagnets and all the sensors to operate over the frequency band 0 to 500 Hz. The need for 4 lateral electromagnets arises because one needs to be able to produce positive forces in any transverse direction.

The actuators had been designed to make the upgrade from a single axis to tri-axial operation possible by incorporating smaller versions of the vertical lift electromagnets on each of the 4 vertical supports of the vertical electromagnet.

The approach adopted to increasing the bandwidth of the sensors was to use pairs of sensors, one effectively a DC sensor and the other an AC sensor, in conjunction with a complementary cross over filter to maximize the output signal-to-noise ratio over the entire frequency band. The sensor combinations chosen were as follows:

(1)   raft motions: a gap sensor for low frequencies and an accelerometer for high frequencies;

(2)   hull or base motions: a gap sensor for low frequencies and an accelerometer for high frequencies;

(3)   force or flux variations: a Hall effect sensor for low frequencies and a multi-turn pick up coil for high frequencies.

These general requirements can be formulated as a set of critical questions as follows.

Critical Questions.

(1)   Can one design a lateral electromagnet, that fits the available space in the single axis actuators, and which has a force capability of 300 Newtons or better.

(2)   Can an integrated Hall effect and inductive coil flux sensor be designed.

(3)   What is the best way of using this integrated sensor to achieve the maximum signal-to-noise ratio over the full frequency band.

(4)   What is the best way of using the inductive gap sensors and accelerometers to provide information to the control system.

(5)   Can the system noise levels be maintained to a satisfactorily low level.

The work carried out to answer these critical questions is described below.

5.1 Sensors

5.1.1 Raft Motion Sensing

To make raft motion measurements over the full frequency band, with maximum signal-to-noise performance, a combination of inductive gap sensors, for low frequencies, and accelerometers mounted on the raft for the higher frequencies was used. The inductive gap sensors work well at low frequencies but their performance falls off at higher frequencies. At higher frequencies accelerometers give better performance.

The gap sensors use magnetic induction at about 500 kHz to sense the gap. As the three orthogonal gap sensors in each actuator were physically quite close together the operating frequencies were staggered to avoid any possibility of mutual interference. Details are given in AST/ERC/12.2/639/TD.

5.1.2 Base Motion Sensing

The initial plan was to infer the base motion by differencing the gap motion and the raft motion, as measured by the raft accelerometers. This requires the raft accelerometers to work in the same frequency range as the gap sensors down to the lowest frequencies. However, in the actual experiments this was superseded by a combination of the gap sensors and 8 accelerometers mounted directly on the base to extract the 6 rigid body base motions.

5.1.3 Force or Flux Sensing

To make flux measurements over the full frequency band, with maximum signal-to-noise performance, a Hall effect sensor was used for the lowest frequencies and an inductive loop for the higher frequencies. A composite flux sensor was designed and tested, using multi-layer printed circuit technology for the inductive coil with the Hall effect sensor integrated into the printed circuit. The Hall effect sensors were Siemens KSY14 which have a high sensitivity and good linearity. Two of these were mounted in the centre of a 12 turn coil incorporated in a 6 layer printed circuit board. One of the Hall effect sensors was a spare. Details of these sensors are given in AST/ERC/12.2/639/TD.

5.2 Signal Conditioning

The voltage output from the individual sensors needs to be centred on the range of the analogue to digital converters (Adc's), namely ± 0.75 V, over the full frequency range. In addition the voltage output from the digital to analogue converters (DAC's), namely ± 0.75 V, needs to be amplified to match the amplifier input requirements. It is important to ensure, in all cases, that there is no possibility of damage to the Adc's or DAC's by any accidental over-voltages being generated.

The signal conditioners provide the following:

(1)   the required gain and off-set;

(2)   filters to reduce out-of-band noise;

(3)   a quick blow fuse in series with a pair of back-to-back "Tranzorbs", type SA5, to provide protection to the Adc's and DAC's.

Earlier signal conditioning units could provide variable gains and variable filters but later units worked with fixed gains and fixed filters. For details see AST/ERC/12.3/630/TD.

5.3 Amplifiers

The control system requires that a linear relationship can be established between the control "demand" and the applied force, which is proportional to the square of the measured flux. The applied force is generated by the current in the electromagnets so this current must be controlled in such a way as to match the "demand". Since the flux generated by an electromagnet is a function of both the current and the gap, any variation in the gap, due to external disturbances, will alter the flux and hence the applied force so the control system needs to be able to cope with these situations. Without any control the electromagnet acts as a very strong negative spring - the force increases with decreasing gap. A high gain analogue flux feedback loop in the amplifier was used to stabilize the flux variations with gap changes to largely overcome this negative spring effect.

At the lowest sea-way frequencies the magnet gaps are controlled to be very nearly constant so the electromagnet acts primarily as a fixed inductance. However, at the highest frequencies the flux is controlled to be very nearly constant, independent of gap variations, so the electromagnet will not be generating any significant back EMF's and thus acts as a pure resistance. In the intermediate frequency range the electromagnet will exhibit similar dynamics to an electric motor since it will be used to oppose the motions of the raft. Thus the means for controlling the current has to cope with a widely different set of loads. There is also the overriding requirement that the control system must have a very low noise output since this noise output will set the limit on the overall capability of the control system.

The results from Phase 1 demonstrated that while the PWM amplifiers were entirely satisfactory for the control system between 10 and 40 Hz, the noise introduced from the PWM amplifiers was far too great for any useful demonstrations of the control system above 40 Hz. It was thus of great importance that a means was found for reducing this noise level to an acceptable value. This led to a short separate investigation of the best means for driving the electromagnets.

Two main types of power amplifier were considered:

(1)   pulse width modulation (PWM) amplifiers, which act as a constant current supply to the electromagnets for a fixed input. This characteristic makes the flux very sensitive to gap variations.

(2)   linear amplifiers which, for a fixed input, act as a constant voltage supply to the electromagnets. This characteristic makes the flux relatively insensitive to gap variations

Show in Table 5.3.1 below is a general comparison between performance that might be expected from these two different approaches.

Feature.	PWM Amplifiers.	Linear Amplifiers.
Flux variations with gap.	Sensitive.	Insensitive.
Low frequency drive.	Very good.	Limited.
High frequency noise output.	Poor.	Good.
Power Efficiency.	Good.	Poor.

Table 5.3.1 Comparison between PWM and Linear Amplifiers

5.3.1 The Experimental Programme to Resolve the Amplifier Choices

The noise level of the PWM amplifiers was measured and found to be -64 dB - relative to 1 volt RMS - over the band 10 to 1,000 Hz. Most of this noise was traced to the current transformer used. This was a Honeywell CSLA2CD. This was replaced by an improved current transformer, a Heme LTA50P/SP1, and the noise immediately fell to -97 dB. This reduction was judged to be sufficient for driving the main lift electromagnets in this programme. All the vertical lift magnets were powered from these modified PWM amplifiers. For details see AST/ERC/12.4/684/EN.

A commercial linear amplifier was tested on the lower powered lateral electromagnets. Their noise levels were very good, -90 dB - but their slew rates were unsatisfactory. By making a relatively simple modification the slew rate was increased to an acceptable level. In the interests of speed and costs, it was decided to power all the lateral electromagnets were from these modified commercial linear amplifiers. The details are described in AST/ERC/12.4/692/EN.

5.4 Lateral Electromagnets

The lateral electromagnets were designed to convert the single axis actuators to tri-axial actuators. These lateral electromagnets were designed to the same patten as the vertical lift electromagnet. The maximum external dimensions are approximately 60% of those of the vertical electromagnet and the design aim was to achieve a steady force, at a gap of 5 mm, of 200 Newtons with a transient capability of 300 Newtons. The corresponding figures for the vertical lift electromagnet are a steady force, at a gap of 5 mm, of 1,200 Newtons. The coil impedance of the lateral electromagnets was designed to be 12 Ohms as compared with 4 Ohms for the axial electromagnet. The coil winding was of copper with 1,560 turns. For more details of these electromagnets see AST/ERC/12.0/639/TD.

5.5 Additional Manufacturing for Raft Upgrade

5.5.1 Flux Sensing

The 16 actuators were manufactured along with 16 vertical integrated flux sensors and 64 lateral integrated flux sensors. These were fitted into the actuators. The lateral AC flux sensors were further integrated as opposing pairs reducing the number of outputs from 4 to 3 per pair.

5.5.2 Gap Sensing

The 48 transverse inductive gap sensors and 48 transverse DC accelerometers were fitted into the actuators.

5.5.3 Signal Conditioning

New, and simpler, signal conditioning units were manufactured to deal with all the additional sensors.

5.5.4 Amplifiers

The existing 16 vertical PWM amplifiers were upgraded and 32 commercial linear amplifiers were purchased, modified as indicated above, and used to drive the lateral electromagnets.

5.5.5 Lateral Magnets

The 16 actuators required the manufacture of 64 new lateral magnets. The lateral magnets were then fitted into the actuators to give them a tri-axial capability.

5.5.6 Commissioning the New Raft

This work involved the following:

(1)   the uprating of the 16 existing actuators to bring them up to the full tri-axial configuration with all their sensors incorporated;

(2)   the testing and calibration of these new uprated actuators before their inclusion into the raft;

(3)   reassembling these into the raft structure;

(4)   connecting up the sensors to the signal conditioning units and the electromagnet drive units to the amplifiers;

(5)   connecting the signal conditioning units to the DSP inputs and connecting the DSP outputs to the electromagnet drive units;

(6)   achieving an initial lift with a very simple point-wise control algorithm.

5.5.7 Summary of Raft Upgrade Work

Reviewing the critical questions listed earlier in this section it is evident that satisfactory lateral magnets were designed, that integrated Hall effect and inductive coil flux sensors were designed and that the system noise levels were reduced to a satisfactory level for the experimental programme to continue.

The only outstanding questions to be answered, at this stage of the work, were the detailed design of the complementary cross-over filters for the gap and accelerometer sensors and for the Hall effect and the inductive coil sensors.

These filter designs were not seen as critical and were deferred until the experimental programme required them.

6 The DSP Upgrade

In the Phase 1 programme 2 B-model DSP cards plus 2 B-model I/O cards had been designed and made. Each DSP card contained an I/O section and 4 Motorola 96,002 DSP processors for the control calculations. The I/O card had just the I/O section. The I/O section could handle 16 analogue inputs and 8 analogue outputs. It included a Motorola 96,002 DSP processor for carrying out the associated I/O functions of decimation, interpolation, anti-alias filtering etc. The general layout of one of these cards is shown in Figure 6.1.
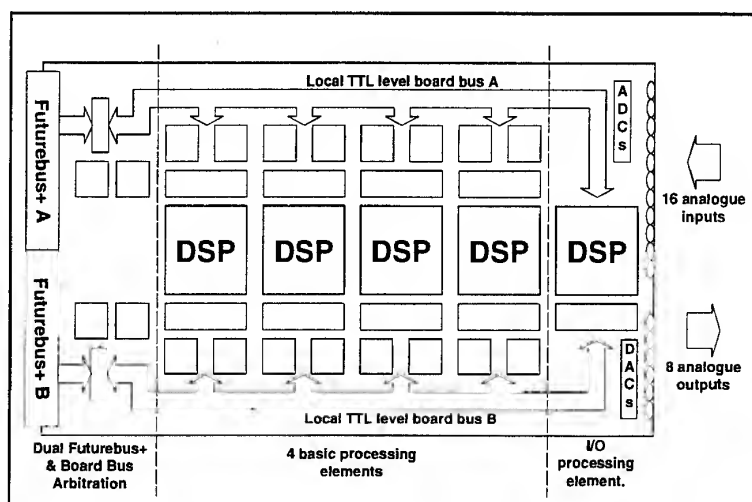

Figure 6.1: DSP Card Layout

Up to twelve of these cards could be connected together via a Futurebus+ backplane and connected to a host PC via a specially designed interface communications card as shown in Figure 6.2.

All the communications between cards, or between a card and the host PC, or within a single card were designed to be asynchronous. The fundamental design rule required that once a communication path had been granted, no other communications could take place until this particular transaction had been completed - other requests would be held until they could be dealt with.

Towards the end of Phase 1, problems had arisen with the logic that controlled the communications between processors. Some of the timings were marginal - busy signals were not always being produced soon enough - and new transactions could begin before a previous transaction was completed. This not only caused occasional corruption of the data, but could, unexpectedly, result in the processor locking up. This possibility was later confirmed by Motorola. During Phase 1 these problems were overcome by a variety of software "fixes" - but this was just not a satisfactory basis for the current programme.

The decision was made to redesign the communications and its associated arbitration logic - some 6% of a fully populated DSP card. The new design made use of faster programmable logic devices, which had become available since the original design, and switched to a fully synchronous logic within a single card - asynchronous logic was retained for inter-card transactions. These new cards were designated C-model cards. This new communications arrangement is illustrated in Figure 6.3.
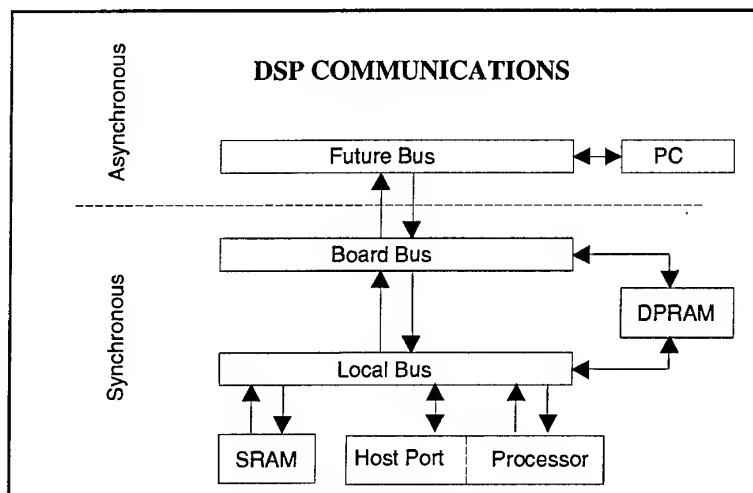
Figure 6.2: DSP Rack with Backplane

Figure 6.3: DSP Processor Communications

An important additional requirement was introduced that not only could all component design tolerances be accommodated but also and that the worst case situation would result in a slip of no more than one clock cycle. A subsidiary decision was to incorporate all the arbitration logic on small "daughter" cards so that if any further logic changes should be required at any time in the future, these could be readily implemented without the need to alter the layout on the main card. The change in the logic on the main DSP cards required that the interface communications card, for connecting a rack of cards to the host PC also needed to be modified.

The main layout of the DSP cards is described in the documents AST/MRC/12/280/TD and AST/MRC/8.0/353/TD.

In addition an inter-rack management card was designed to handle communications between racks of DSP cards or between a rack of cards and the host PC. This would enable more than one rack of cards to be used in the experimental work if this proved to be necessary.

6.1 DSP Hardware Design, Testing and Manufacture

To implement the new C-model design, the tasks undertaken are listed below.

(1)   The re-design of all the communications logic between individual processors and the arbitration logic required.

(2)   The re-design of the interface communications card, for connecting a rack of cards to the host PC.

(3)   The design of an inter-rack management card.

During the design process a number of detailed design reviews were carried out before moving onto the next stage. Details are to be found in the documents AST/MRC/8.0/615/EN and AST/MRC/12/674/TR.

During Phase 1 problems had arisen with the manufacture of both the bare PCB cards and with the incorporation of the electronic components. A decision was made that, in this programme, the bare card manufacture would be done by one sub-contractor, the card connectivity 100% tested by a second sub-contractor and the final assembly done by a third sub-contractor. The final acceptance tests were carried out by technicians working with the original designers. The manufacturing plans are listed in AST/MRC/12.5/672/TR.

The next stage was the manufacture of two new C-model cards and the testing of all the new communications logic.

This consisted of the following steps.

(1)   The layout of the new PCB cards.

(2)   The manufacture and testing of new PCB cards.

(3)   The construction of a two new DSP cards.

(4)   The testing of all the new communications logic incorporated.

(5)   The testing of the Futurebus+ backplane which interconnects cards within a single rack.

When this stage had been partially completed the manufacture of the remaining cards likely to be required in this programme was started to avoid any further delays. Further, more detailed testing of the new cards continued and one or two minor logic corrections were made in the programmable logic devices.

The final raft ultimately required 5 outputs per actuator, one for each electromagnet, and up to 16 inputs per actuator made up from 10 flux sensor inputs, 3 gap sensor inputs and 3 DC accelerometer inputs. However, in the interests of economy, 2 pairs of opposing lateral flux sensors were connected in opposition in their signal conditioning units, thus reducing the number of sensor inputs per actuator to 14. In addition there needs to be some additional input capability for purely monitoring purposes. The planned manufacture was for a total of 20 cards, 12 of which were full DSP cards and 8 were I/O cards.

6.1.1 Summary of DSP Upgrade

A critical question for Phase 2 was can the new C-model DSP card be designed so as to eliminate all the communications and marginal timing problems that were encountered with the B-model system. The results have show that the new C-model cards do indeed work correctly and that their reliability is excellent. This has enabled the experimental programme to proceed smoothly. By the October 1993 review, the new electronics was made, installed and was working very well indeed although, at this stage, no processor to processor communications between different cards had been required.

As the programme proceeded to its last stages, during November and December 1993, inter-board communications were required and these were tested very thoroughly prior to their use. Communications between individual processors worked very well but very occasional problems were encountered which resulted in data corruption. This took a considerable time to identify correctly. The problem was identified as being associated with the special global write transaction. This global write is used to rapidly transfer all the input sensor data to all the DSP processors in the system. If a single processor to processor transaction was requested during a global write transaction this should be delayed until it could be dealt with. This worked perfectly if the transaction was within a single card but, when the transaction was between cards, the request was very occasionally never completed.

The problem has been identified and it can be corrected but this will require a modification to the daughter cards which control the arbitration logic. Unfortunately, when this problem was identified, there was not sufficient time to correct it before the end of Phase 2. The decision was therefore taken that, during the remainder of Phase 2, the use of the global write instruction would be avoided and all the input sensor information would be passed by a series of single writes. The consequence of this was that passing of even limited sensor information took up some 50% of the total cycle time and put a corresponding limit on the control calculations. Nevertheless, even with this limitation, the low frequency control algorithms were demonstrated with little difficulty although it was not possible to use more than 8 of the actuators in their full triaxial mode. The limitation on the high frequency control algorithm was more severe and its use limited to the lift electromagnet on one actuator only.

6.2 DSP Software

The control algorithms are developed in "MATLAB" - proprietary package for use on a PC. This package is particularly useful for control computations. The basic "units" of information are vectors and matrices at the "MATLAB" level.

When the algorithms are completed and tested they can be output, by the PC, as a series of commands in 'C'. The purpose of the Virtual Machine Package is to make the DSP operate in a way very similar to "MATLAB" but making the most efficient use of the DSP chip's efficient pipe-lining operations along with the parallel processing capability built into the DSP system. Parallel processing can take two forms:

(1)   independent calculations can be carried out in parallel on different sets of chips;

(2)   a particular calculation can be partitioned so that the separate parts of the calculation can be conducted, in parallel, on different chips and the results reassembled at the end of the calculation.

The other requirements for the Virtual Machine Package are:

(1)   to make the coding simple and easy to implement or change;

(2)   to handle all the I/O operations;

(3)   to handle all the communications to and from the host PC.

6.2.1 The Basic Structure

To use the DSP chips efficiently one must use optimized codes written in Assembler - any attempt to use 'C' will slow the system down to about 10% efficiency. Thus at the lowest level the bulk of the processing must be carried out with optimized Assembler coded macros. At the highest level one uses compiled 'C' coded routines to control the use of the Assembler routines. These basic ideas are implemented as follows:

(1)   one DSP processor is designated as the Master Processor. It runs the compiled 'C' coded routines;

(2)   the Master Processor then has to control the use of a number of Assembler coded macros which are stored in the DSP cards;

(3)   the Assembler coded macros are written, as far as practicable, to have a one-to-one correspondence with the commands in "MATLAB".

(4)   these Assembler coded macros are part of a library that can be accessed by the compiled 'C' code in the Master Processor;

(5)   the macros are run on Processors other than the Master Processor.

The initial specification is contained in AST/MRC/8.0/425/TD and further information in AST/MRC/13/673/TR. The virtual machine software and the library of assembler coded macros has been updated, whenever necessary, in order to carry out the experiments efficiently.

7 The Experimental Programme and the Results

The experimental programme set out to answer the following set of critical questions.

Critical Questions.

(1)   How can one gain a sufficiently stable lift so that system identification tests can be carried out.

(2)   How accurately can the modal parameters be measured and how is this information to be fed to the control algorithms.

(3)   How well and over what frequency band can the raft be supported by a simple modal isolation strategy.

(4)   How will one make the transition from the low frequency modal description to progressively more continuous and more localised higher frequency disturbances when natural damping begins to reduce the coherence length to dimensions less that those of the raft.

(5)   When is it necessary to change from feedback to adaptive feedforward control and how can one effect a smooth transition from one to the other.

(6)   How well can the basic mount resonances be suppressed.

(7)   How well does the sea-way motion control strategy work.

(8)   Can one suppress the longer wavelength bending resonances on the raft using shorter wavelength force distributions.

(9)   To how high a frequency should one aim to suppress internal structural resonances in this way.

(10)   How will one make a smooth transition from this regime to a modal isolation regime.

(11)   Are there any dynamical or numerical stability problems associated with increasing the size of the raft.

(12)   Does mass loading introduce any major problems for the control system.

The principal tasks undertaken in the experimental programme are described below.

7.1 Stable Lift and System Identification Tests

The early system identification tests were undertaken on the original 16 vertical actuator raft shown in Figure 3.1.1 since it was not until August 1993 that the necessary manufacture had been completed ant the raft stripped down and converted to the final raft version in Figure 3.2.1. To achieve the initial stable lift the control system was designed to make each actuator independently mimic a spring and damper. The basic algorithm is illustrated in Figure 7.1.1. The 16 gap sensors are processed through a stage where any residual off-sets were subtracted out and normalizing gains are introduced before the sensor data is passed on to 16 independent IIR filters. These generate the required spring and damper characteristics. Also shown is the means whereby additional disturbance signals can be injected to the actuators for test purposes.

Although this point-wise control of the actuators is simple in principle, in practice, it was quite a severe test on the accuracy of the equalisation filters and on the analogue flux feedback control which try to maintain the linearity of the output force as a function of the input demand, independently of gap variations, over the entire frequency band. Errors can easily result in the unwanted excitation of higher frequency resonant modes of the raft.
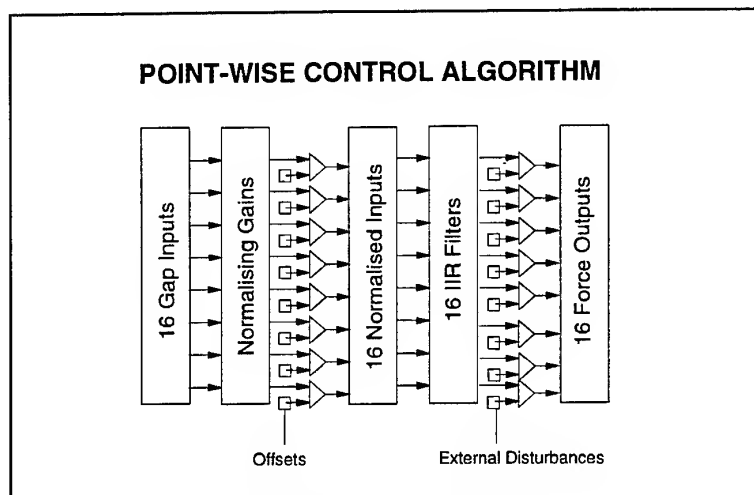
POINT-WISE CONTROL ALGORITHM

Offsets

External Disturbances

Figure 7.1.1: Point-Wise Control Algorithm (Phase 1)

Once this initial stable lift was achieved, random external disturbance signals, generated by the host PC, were injected into each of the 16 actuator electromagnets in turn with the remaining actuator external disturbance inputs clamped. As each actuator was excited, sets of 10,000 data samples were captured from the gap sensor, the raft accelerometer and the DSP output signals. Each data set was then transformed, off line, into a 1,024 sample transfer function. This provided a complete data base of 16 input, 16 output transfer functions for the response of the accelerometers and gap sensors to the individual control signal inputs.

When the raft had been converted to 16 triaxial mounts there were now 48 control forces to generate. The basic initial lift algorithm was modified as shown in Figure 7.1.2. The 48 gap sensors were processed through the stage where any residual off-sets and normalizing gains are introduced before the sensor data is passed on to 48 independent IIR filters which generate the required spring and damper characteristics. Since the lateral electromagnets in the actuators are at an angle of 45° to the X and Y axes the outputs from the inverse transform have to be further processed by the rotation matrix as shown in Figure 7.1.2. Also shown is the means whereby additional disturbance signals can be injected to the actuators for test purposes.

The final stage where the force demand is converted into the amplifier drive voltage is shown in Figure 7.1.3. Opposing pairs of lateral electromagnets were driven in a push-pull manner with the bias current and the generation of the two individual opposing drives being generated in the output conditioning unit.

Although these more complex algorithms were not too difficult to put into place there were, at this stage, problems with having sufficient DSP boards available to make use of the very greatly increased number of sensors - the increase was from 32 to 256. At this stage we had to use the minimum number of sensors required to demonstrate key results. The basic modes we wished to control in the lateral direction were the rigid body modes and these could be excited with 4 to 8 actuators with full lateral control. This was the maximum number we could control with the limited DSP capability then available.
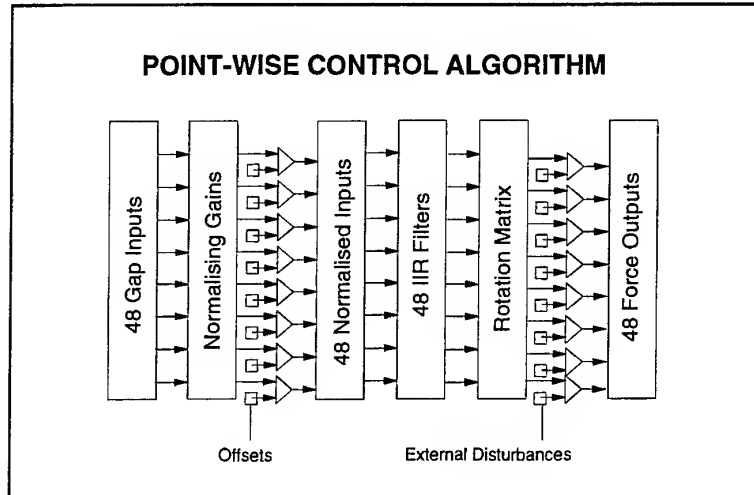
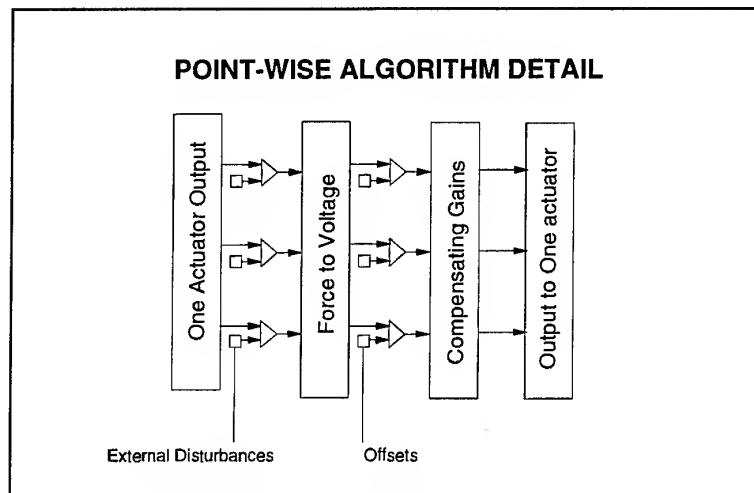**Figure 7.1.2: Point-Wise Control Algorithm (Phase 2)**

Figure 7.1.3: Point-Wise Algorithm Detail

The basic technique for determining the structural modes of the raft was to inject random signals into each actuator in turn and measure the accelerometer and gap response at all the other actuators in turn. Normally 10,000 data points were captured for each sensor along with the excitation data which produced the disturbance. Each data set was then transformed into a 1,024 point frequency spectrum. This data allowed one to determine the transfer functions between the excitation data and the sensor response.

7.1.1 Results and Discussion

Once the raft was lifted test disturbance could be injected at the actuators and the raft response at a point measured. The frequency spectrum of these results are shown in Figure 7.1.1.1. It will be seen that there is a mount resonance at just over 3 Hz and the structural resonances then occur in pairs. The first pair have resonances at 15 and 20 Hz.


Figure 7.1.1.1: The "Passive" Response of the Raft

The modes of the bare raft split into symmetric and anti-symmetric modes - the symmetry being about the centre line. The 6 rigid body modes comprise the symmetric modes: heave; surge; and sway and the anti-symmetric modes comprise: roll; pitch; and yaw. Preliminary calculations indicate that the lowest frequency structural mode is the anti-symmetric torsional vibration mode with the next higher frequency modes being the symmetric and the anti-symmetric U-modes.

The object of these system identification measurements is to firstly determine the "raw" transfer functions - ie where the input force is applied to the actuator under test and all the response at other actuators are measured. These "raw" transfer functions contain the effects of the low frequency control algorithms used to support the raft in the first place. The next step is to deconvolve the effects of the control algorithms to obtain the "free-free" transfer functions. This is done as follows:

(1)   for a specific set of external disturbance inputs, one derives a set of control system transfer functions, $u_{i,j}$, gap transfer functions, $d_{i,j}$, a set of accelerometer transfer functions, $a_{i,j}$;

(2)   construct the "raw" transfer function matrices D (ω_n) A(ω_n) and U(ω_n) from the gap, accelerometer and DSP output transfer functions, where the row index is the actuator at which the measurement was made and the column index is the actuator at which the external disturbance input was applied and n is the index for the particular frequency from the set of 1,024 frequencies;

(3)   the "free" gap and acceleration transfer function matrices, T_d(ω_n) and T_a(ω_n) are then defined by the equations:

$$D(\omega_n) \quad = \quad T_d(\omega_n) \quad U(\omega_n) \qquad\qquad 7.1.1.1$$

and

$$A(\omega_n) \quad = \quad T_a(\omega_n) \quad U(\omega_n) \qquad\qquad 7.1.1.2$$

Thus it follows that:

$$T_d(\omega_n) \quad = \quad inv\,[\,U(\omega_n)\,] \quad D(\omega_n) \qquad\qquad 7.1.1.3$$

and

$$T_a(\omega_n) \quad = \quad inv\,[\,U(\omega_n)\,] \quad A(\omega_n) \qquad\qquad 7.1.1.4$$

These 1,024 matrices $T_d(\omega_n)$ and $T_a(\omega_n)$, corresponding to the different frequencies in the frequency domain, contain all the dynamical information about the "free" raft. Figure 7.1.1.2 shows an example of a "raw" force to gap transfer function and Figure 7.1.1.3 shows the results of the deconvolution process to obtain the "free" force to gap transfer function. It is immediately apparent that there is a dramatic increase in detail in going from the "raw" to the "free" transfer function. The various lower frequency structural resonances stand out clearly in Figure 7.1.1.3 although above about 40 Hz the resolution in the data is insufficient to pick out the higher frequency structural resonances clearly.

Figure 7.1.1.2: Raw Transfer Function - Gap to Force

The final step in the system identification process is to pick out, from Figure 7.1.1.3, the force to gap transfer function matrix at the frequency of one of the specific resonances - say ω_r. The dominant eigenvalue of this matrix defines the dominant frequency - the resonant frequency - and the corresponding eigenvector defines the dominant mode shape. This transformation can be written as follows:

$$T(\omega_r) \quad = \quad \tilde{V}(\omega_r)\,T_d(\omega_r)\,V(\omega_r) \qquad\qquad 7.1.1.5$$

where $V(\omega_r)$ is the eigenvector matrix. Figures 7.1.1.4 to 7.1.1.9 show the mode shapes of the lower frequency modes determined in this way.

Figure 7.1.1.3: Free-Free Transfer Function - Gap to Force

As a final cross-check the dominant eigenvalue $T(\omega_r)$ is the corresponding modal transfer function. These modal transfer functions should contain little, if any, evidence of other modes. Figures 7.1.1.10 to 7.1.1.14 show the modal transfer functions determined in this way and it will be seen that they do not contain mixtures of other modes - Figure 7.1.1.14 is a good example.

It will be seen that the torsion mode has a resonance at just over 8 Hz and the symmetric U-mode has a resonance at about 15 Hz. The absence of other resonances in these figures indicate that the dominant mode is, in fact, very dominant and thus the system identification technique is accurate. If further precision were required one could excite single modes only and measure the transfer functions in this case.

It will be seen that in the case of the heave mode transfer function, Figure 7.1.1.10, that there is evidence of some residual negative spring action at the lowest frequency. This is due to the method of measuring the forces. These were assumed to be defined by the DSP output voltages but one must note that the DSP output voltage is subject to further modification by the local analogue flux feedback control. The effect of this local analogue flux feedback control is to reduce the strong "negative spring" action of a normal electromagnet to a "weak spring" action by making the flux very nearly independent of gap variations.

In summary the system identification techniques described above can be conducted with the raft levitated, as indicated. The techniques are relatively straightforward and the results obtained indicate that the lower frequency modes can be identified accurately and the mode shapes agree with the predicted mode shapes for this simple structure.

Heave Mode

Normalised Displacement

0

Actuator Position

Front Row
Back Row

Figure 7.1.1.4: Heave Mode Shape

Pitch Mode

Normalised Displacement

0

Actuator Position

Front Row
Back Row

Figure 7.1.1.5: Pitch Mode Shape

Figure 7.1.1.6: Roll Mode Shape

Figure 7.1.1.7: Torsion Mode Shape

Symmetric U-Mode

Normalised Displacement

Front Row
Back Row

Actuator Position

Figure 7.1.1.8: Symmetric U-Mode Mode Shape

Anti-Symmetric U-Mode

Normalised Displacement

Front Row
Back Row

Actuator Position

Figure 7.1.1.9: Anti-Symmetric U-Mode Mode Shape

Figure 7.1.1.10: Heave Mode Transfer Function

**Figure 7.1.1.11: Pitch Mode Transfer Function**

Figure 7.1.1.12: Roll Mode Transfer Function

**Figure 7.1.1.13: Torsion Mode Transfer Function**

Figure 7.1.1.14: Symmetric U-Mode Transfer Function

7.2 Low Frequency Modal Isolation

Once the modal structure of the raft had been determined, as described in section 7.1, it became possible to control the raft using modal control. This involves the transformation of the individual gap inputs to modal gaps using the eigenvector transformation matrix. The central part of this modal control technique is illustrated in Figure 7.2.1. The actual algorithm structure used for the modal control experiments is considerably more complex and is illustrated in Figure 7.2.2.

The input to the modal transformation matrix consisted of the 48 gaps. The 24 outputs from the modal transformation matrix were chosen to be the 6 rigid body modes and 18 structural resonant modes. There are, of course, many more than 18 modes but if the eigenvector of a particular mode is not included in the modal transform matrix, it will not be controlled and will simply vibrate in an unconstrained manner. Such a freely vibrating mode cannot produce any net force transmission through the actuators and on to the base.

Each of these 24 modes could be controlled separately by means of the corresponding gain and IIR filter as shown. The outputs from these filters are then passed on to the inverse modal transform matrix to generate the force demands for the 16 tri-axial actuators. The processing of these force demands through the rotation matrix and on to the drive amplifiers is as shown in Figure 7.1.2 and Figure 7.1.3.
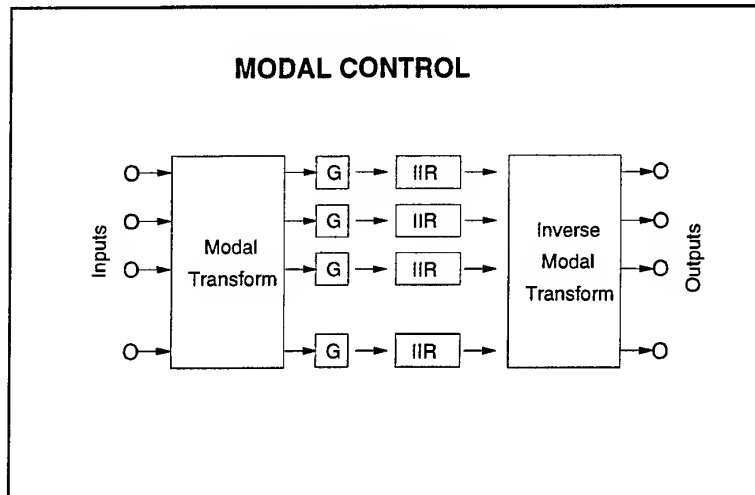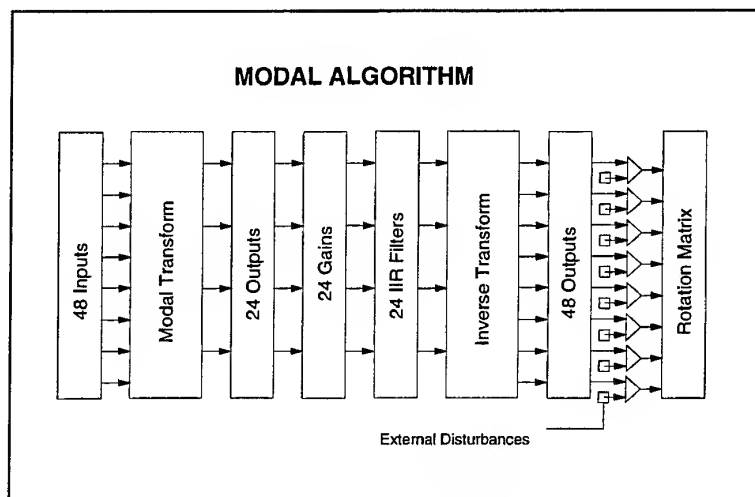
Figure 7.2.1: Modal Control Algorithm - Outline

Figure 7.2.2: Modal Control Algorithm - Detail

7.2.1 Results and Discussion

Figure 7.2.1.1 shows the results of supporting the raft with modal control and exciting the raft at a point. The first curve shows the result with no damping and the second with damping applied selectively to the symmetric and anti-symmetric U-modes. The modes above 40 Hz were allowed to run free. The mount resonance has been virtually eliminated by damping all the rigid body modes.

Figure 7.2.1.1: Modal Response to Point Excitation

A further demonstration of the low frequency modal control is obtained when one excites the raft with a single mode force pattern. If the modal identification is correct this should excite the U-mode only - any excitation of other modes would indicate that the shape is not accurately orthogonal to these modes. The results of doing this with the symmetric U-mode - ie a U-force - is shown in Figures 7.2.1.2 and 7.2.1.3.


Figure 7.2.1.2: Single Mode Excitation - Modal Response

It will be seen that the excitation does, in fact, generate a response for the single U-mode only. There is, of course, residual excitation at frequencies above 40 Hz. This is the region where modes are allowed to run free and where the spatial resolution of the U-force shape is no longer able to be accurately orthogonal to these short wavelength resonances.

Figure 7.2.1.3: Single Mode Excitation - Force Transmission

In addition Figure 7.2.1.2 shows the ease with which the symmetric U-mode can be selectively and progressively damped. The effect of this progressive damping on spreading the frequency band over which the force is transmitted to the base is clearly seen in Figure 7.2.1.3.

These results illustrate the flexibility of modal control in that it allows one to damp out mount resonances, apply damping selectively to some modes while allowing other modes to run free. This latter feature is especially important since these higher frequency shorter wavelength modes are not sufficiently well spatially sampled, by the raft sensors and actuators, to make their control physically possible. However, it is easy to allow them to run free knowing that they will produce no net force into the base.

7.3 Sea-way Motion Control

The original intention for the sea-way control had been to derive the base acceleration from the differences between the gap acceleration and the raft acceleration. This was the scheme used in the control theory simulations - see section 4.2.2. Instead, however, a simpler feedforward scheme was implemented where the base accelerations were measured directly by means of 8 accelerometers mounted as shown in Figure 7.3.1.

The outputs from these 8 base accelerometers were then transformed into the 6 rigid body modes of the base, low pass filtered, with a cut-off at about 1.5 Hz, and then used to add the additional forces required to the raft rigid body modes to track the base movements. The principle is illustrated in Figure 7.3.2 and the complete algorithm structure is shown in Figure 7.3.3.

This control strategy made the raft seem very stiffly mounted to base disturbances for very low frequency disturbances, such as simulated sea-way motions, while remaining very softly mounted for vibrations originating on the raft. The entire table supporting the raft was moved hydraulically to simulate the accelerations encountered in sea-way motions.

Figure 7.3.1: Base Accelerometer Layout
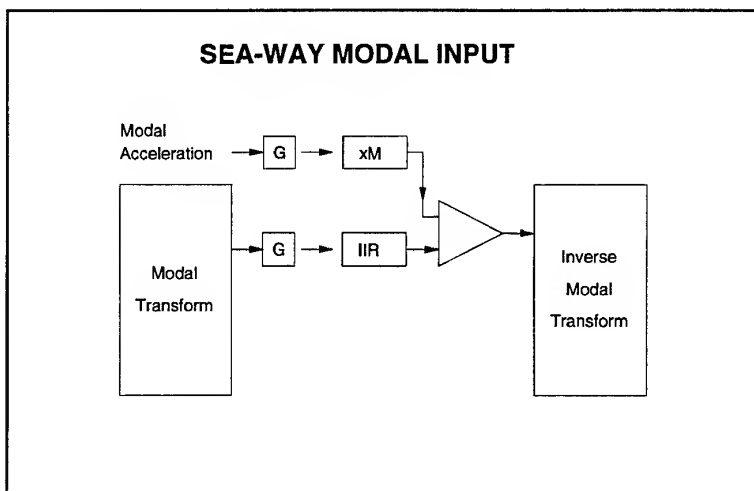
Figure 7.3.2: Seaway Input to Modal Algorithm

MODAL TRANSFORM WITH BASE MODES ADDED

Figure 7.3.3: Seaway Input to Modal Algorithm - Detail

7.3.1 Results and Discussion

The supporting base was driven hydraulically with a period of 10 seconds and the lateral motion of the raft was measured firstly with no sea-way control added and then with the addition of the sea-way motion control. Figure 7.3.1.1 shows the time history response over a 20 second period.

Figure 7.3.1.1: Suppression of Seaway Motion

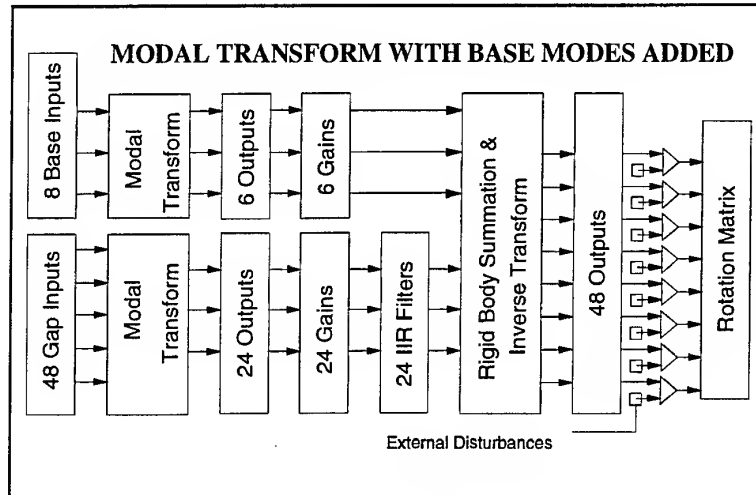It will immediately be seen that in the absence of the sea-way control there is a large lateral surge displacement of the raft. In fact the raft is hitting its end stops or the measured surge would be even greater. The transient that occurs on the positive displacement arises from the control mechanism of the hydraulic ram. When the sea-way control is included there is a dramatic decrease in this surge response as shown by the second curve in Figure 7.3.1.1. The surge displacement is reduced to about 0.1 mm.

These results can be expressed more quantitatively by determining the power spectral density of the sea-way input, as measured by the 8 base accelerometers, and the power spectral density of the raft response both with and without the sea-way control. The first of these results is shown in Figure 7.3.1.2 where it is clear that the control action has virtually no effect on the base accelerations. The second result, shown in Figure 7.3.1.3, indicates that at 0.2 Hz there is a 30 dB reduction in the power spectral density of the raft response. The lowest frequency for these results, 0.2 Hz, was set by the length of the records we could collect from a single experiment. Figure 7.3.1.3 also shows that this sea-way control is rolled off at 1.5 Hz.

The next demonstration was to show that the raft response is unaffected by the addition of the sea-way control. Figure 7.3.1.4 shows the time history of the lateral displacement of the raft to an impulsive lateral disturbance. Here it is clearly seen that the response is the same both with and without the sea-way control activated showing that the raft remains soft to disturbances generated on it.


Figure 7.3.1.2: Power Spectral Density of Seaway Input

Power Spectral Density of Raft Motion

Magnitude (dB) vs Frequency (Hz)

o - Without Control
+ - With Control

Figure 7.3.1.3: Power Spectral Density of Raft Response

Vibration Response with & without Seaway Control

Surge Displacement (mm) vs Time (seconds)

Figure 7.3.1.4: Vibration Response with & without Seaway Control

The results of a more complex demonstration are shown in Figure 7.3.1.5. Here the raft is directly subject to an external disturbance while its base is being driven by the hydraulic actuator with a period of about 10 seconds. It will be seen that in the absence of the sea-way control the external disturbance is superimposed on the response of the raft to the movement of its base - see Figure 7.3.1.1. When the sea-way control is activated this same disturbance is seen to be superimposed on the raft as it now tracks the movement of its base.

Figure 7.3.1.5: Lateral Response to Seaway & Vibration Inputs

In summary these demonstrations indicate that with the sea-way control in operation, the raft is stiff to low frequency inputs from the moving base while remaining soft to disturbances generated directly on the raft. The magnitude of this non-reciprocal behaviour is seen, from Figure 7.3.1.3, to be up to 30 dB.

7.4 Structural Resonance Suppression and Isolation

The results illustrated in Figure 7.2.1.1 show that structural resonances can be controlled by simple damping. However, simple damping can be augmented by generating additional forces that oppose the vibration of a resonant mode - in other words to "stiffen" this particular mode. This is particularly easy for the longer wavelength bending resonances and the simplest method is to use an opposing force pattern with the same shape as the mode one is wishing to control. By applying this force pattern with the correct strength, but in anti-phase to the resonant mode, one can suppress this particular mode and thus "stiffen" the raft.

In the case of these longer wavelength bending resonances there is no need to be restricted to making the control force pattern have the same shape as the mode being controlled. By using shorter wavelength force patterns it should not only be possible to control the amplitude of the longer wavelength modes on the raft but to do so in a way that simultaneously reduces the efficiency of any coupling into propagating acoustic modes in the sea. This topic was dealt with in section 4.2.4.

At this stage of the programme the DSP system was far to small to handle the full LQG calculations described in the section on control theory and simulations above. The approach adopted was to examine, off-line, the control force patterns that could be used to control the symmetric U-mode bending resonance of a beam firstly with a force pattern of the same shape, the U-force shape, and then with an optimal force pattern, derived from the LQG model, to minimize the acoustic radiation, the W-force shape. These two force shapes are shown in Figure 7.4.1 - note the greater amplitude and shorter wavelength of the W-force shape compared to the U-force shape.

Force Patterns to Control the Symmetric U-Mode

(Force Amplitude vs Beam Datum Point)

U-Force Shape -----
W-Force Shape ———

Figure 7.4.1: Theoretical Control Force Shapes

The wavenumber spectrum of these two force shapes are shown in Figure 7.4.2 where it will be seen that at low wavenumbers the spectrum of the more complex W-force shape is some 20 dB lower than that of the simpler U-force shape. It is this feature that has the greatest impact in reducing the acoustic radiation that will arise from the reaction to this control force in the hull. This is shown more clearly in Figure 7.4.3 which shows the attenuation of the low wavenumber spectrum when the W-force shape is used in preference to the U-force shape.

Wavenumber Spectrum for Control Force Shapes

(Amplitude (dB) vs Wavenumber)

U-Force Shape ----
W-Force Shape ———

**Figure 7.4.2: Wavenumber Spectrum for U-force & W-force Shapes**

Figure 7.4.3: Low Wavenumber Attenuation for W-force Shapes

The low frequency modal control, described in section 7.2 was modified so that the control force applied to the symmetric U-mode of the raft was close to the W-force shape and the actual force shapes used are shown in Figure 7.4.4.


Figure 7.4.4: Actual Control Force Shapes

7.4.1 Results and Discussion

The point excitation results, obtained with the U-force control from Figure 7.4.3, are shown in Figure 7.4.1.1. Also shown in Figure 7.4.1.1, for comparison, are the results for the "passive" raft. It will be seen that the U-force shape is able to eliminate the symmetric U-mode resonance - the reduction is over 15 dB.

The point excitation results, obtained with the W-force control from Figure 7.4.3, are shown in Figure 7.4.1.2 and again, for comparison, are the results for the "passive" raft are also shown. It will be seen that the W-force shape is also able to eliminate the symmetric U-mode resonance and again the reduction is over 15 dB.

7.5 High Frequency Localised Isolation

As described in section 4.3 the situation is likely to arise where each mount is subject to a component of unwanted leakage flux variation due to the local vibration levels associated with each magnet armature which, in turn, generates a residual force transmission. The algorithm aims to sense the high frequency motions of a magnet armature and use this information, in a feedforward control, to minimize this residual force transmission. The residual force transmission can be monitored, using the AC flux loops, and this information can be used to make the feedforward control adaptive.

7.5.1 Complementary Cross Over Filters

In order to have a wide bandwidth measure of the magnet gaps for the high frequency work the output from a gap sensor and the corresponding accelerometer were combined, with complementary low and high pass filters, to give their optimum signal-to-noise ratio over the full frequency band 0 to 500 Hz. The output from the accelerometer is doubly integrated to convert the results to gap variations as shown in Figure 7.5.1.1. The optimum cross-over frequency for these devices is 7 Hz and the complementary filter characteristics are shown in Figure 7.5.1.2. The filters were designed so that complex sum of the outputs from the two filters is exactly unity over the full Nyquist bandwidth.


**Figure 7.4.1.1: Control with U-Force Shape**

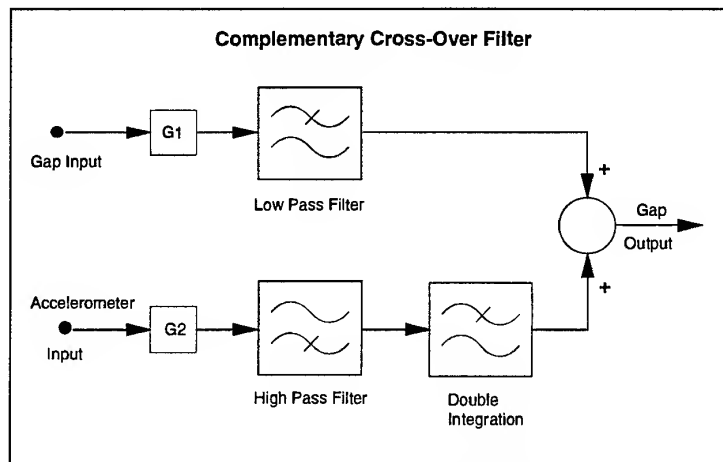**Figure 7.4.1.2: Control with W-Force Shape**

Figure 7.5.1.1: Complementary Cross-Over Filter Arrangement

These cross-over filters were implemented on all the vertical electromagnets although the high frequency algorithm investigations had to be limited to a single actuator because of the limitations of the DSP system.

Figure 7.5.1.2: Cross-Over Filter Characteristics

7.5.2 Initial Correlation Measurements

The first stage of this work was to measure the correlations between the accelerations on one mount and the AC flux loops, on the same mount, for a random disturbances at other points on the raft. The performance of the high frequency algorithm is critically dependent on these correlations being high - if there is no correlation any attempt to control the flux variations by a feedforward control from the accelerations will have no effect.

The first set of measurements showed that the background noise levels were unacceptably high and that good coherence could only be obtained for test signals of a very high level applied directly over the mount.

The major source of noise was from out-of-band noise from the low frequency algorithm. The low frequency algorithm was very substantially improved, by providing increased high-frequency roll-off for the rigid body, torsional and long wavelength bending modes. The low frequency algorithm continued to give optimal damping at the resonant frequencies of these modes and, although the characteristics were empirically defined, they correspond very closely with the optimal LQG diagonal modal control description. This new low frequency algorithm proved to be very stable and extremely quiet - the best performance that has yet been achieved in this programme. The improvement is illustrated in Figure 7.5.2.1.

Once this new low frequency algorithm was in place it was possible to systematically track down and minimize other noise sources. Some of these are listed below:

(1)   the DSP outputs required additional low-pass filtering and, to preserve the overall linearity, a complementary high-pass filter was included within the DSP;

(2)   the sensor inputs to the DSP can have very large dynamic ranges which cannot always be adequately catered for by the fixed, 14 bit resolution of the Adc's. Suitable pre-emphasis/de-emphasis filters need to be introduced, one on each side of the ADC, to keep the dynamic range correctly balanced;

Figure 7.5.2.1: Noise from Low Frequency Algorithms

(3)   the noise output from the PWM amplifiers could be further improved by either reductions in their bandwidth or by further improvements to the amplifier current sensing components - see section 5.3.1.

Although it was not possible, in the time available, to incorporate all these improvements, the overall noise level was reduced by 40 dB.

With this new lower noise floor the correlation measurements were repeated and the transfer functions between acceleration and flux measured for a single actuator. The results are shown in Figure 7.5.2.2.

These measurements show that the mount responded to inputs in non-orthogonal directions, indicating that the control system will require a full cross-coupled matrix description - a 3×3 cross-coupling matrix of transfer functions. In the simulations it had been assumed that the matrix of these correlations was diagonal - ie X accelerations correlated with X flux changes only, with similar results in the Y and Z directions.

Figure 7.5.2.2: Transfer Functions - Acceleration to Flux

7.5.3 Modified Algorithm

As a result of these investigations it was decided to use a simple open-loop feedforward control strategy to test the high frequency algorithm where the appropriate transfer functions were defined off-line as a result of these correlation measurements. The basic scheme is shown in Figure 7.5.3.1 which shows the outputs from the 3 orthogonal accelerations being passed through three filters with variable gains to be subtracted from the output to the vertical electromagnet on one mount - this arrangement can be repeated for all three mount axes. Figure 7.5.3.2 shows how this is coupled with the low-frequency output to this same vertical electromagnet.

7.5.4 Results and Discussion

As will be seen, from the previous description, the high frequency control algorithm was the most difficult to test satisfactorily. The details of the coupling between vibration inputs at one point on the raft to the force transmission at another would be very difficult to simulate accurately before the event. In addition the out-of-band noise from the low frequency algorithm has a major influence on how well the high frequency algorithm can work. This is an important part of the design of the low-frequency algorithm. In the event the programme was terminated just before the high frequency algorithm could be demonstrated although little work remained to be done.
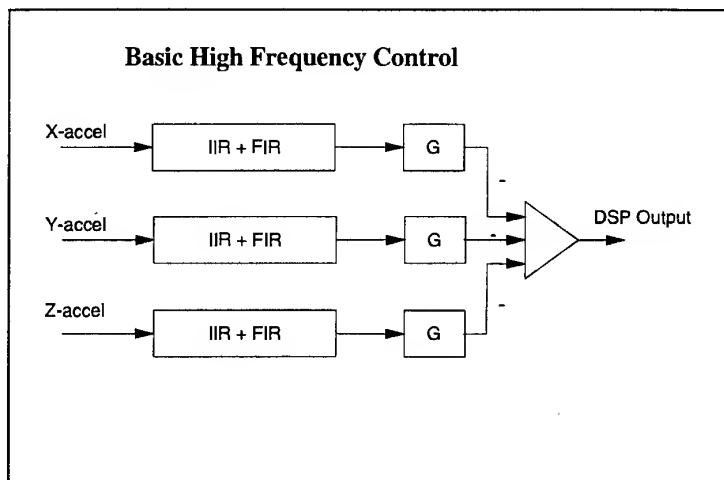
Figure 7.5.3.1: Basic High Frequency Algorithm
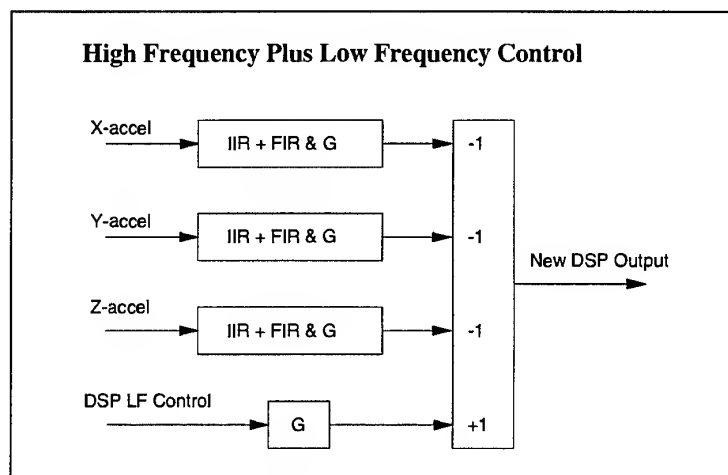
Figure 7.5.3.2: Low & High Frequency Algorithm

The predicted performance, based on the transfer functions shown in Figure 7.5.2.2, are shown in Figures 7.5.4.1 and 7.5.4.2. The first of these shows the performance based on using the vertical transfer function data only and the second the effect of using all three transfer functions. It will be seen that, as expected, the triaxial control predictions are considerably better than the single axis control. The predicted improvements range from 15 to 35 dB.

Figure 7.5.4.1: HF Performance - Vertical Control Only

Figure 7.5.4.2: HF Performance - Tri-axial Control

It will also be seen that these results do not predict any significant performance improvements above 150 Hz. This is because the algorithm is based on fixed transfer functions from any arbitrary point to the actuator being controlled. This is in line with the results of the correlation measurements which indicate that the correlations become specific to particular paths above 150 Hz. However, with the use of adaptive control techniques, the performance improvement will be optimized for a specific set of high frequency excitations arising from fixed machinery so one would expect the performance, in such cases, to extend to much higher frequencies.

7.6 Stability Problems with Increasing Size

Finally there is the question: are there dynamical or numerical stability problems, associated with increasing the size of the raft and, if so, how can they be overcome?

The answer to this question was clarified during the experimental work. In the initial high frequency measurements of correlations - section 7.5.2 - one found that the dynamic range of the Adc's was insufficient and precision was being lost. This problem can be cured, as indicated, by the use of suitable pre-emphasis/de-emphasis filters, one on each side of the ADC, to keep the dynamic range correctly balanced. No other significant numerical problem was encountered.

The dynamical stability of the control system was always good. The only occasions where there were stability problems - apart from coding errors - were due to the residual "negative spring" effect produced by the high gain analogue flux feedback control. This control aims to make the flux - and hence force - produced by an electromagnet independent of variations in the gap. As the loop gain is finite, there is inevitably a residual "negative spring" effect. In most cases, the main control system had no difficulty in overcoming this problem. The use of this flux control was always seen as a temporary solution which would no longer be needed when the DSP system was able to take over the total control of the electromagnets.

To summarise, the work to date has not shown up any significant numerical or dynamical stability problems that have not been easily dealt with. This gives one considerable confidence that any such problems will be as easily dealt with in the future.

7.7 Preliminary Effects of Mass Loading

To investigate some of the preliminary effects of mass loading a 250 kg mass was mounted on the centre line of the raft at one end so that this additional mass was largely supported on the 4 end actuators - see Figure 7.7.1.


Figure 7.7.1: Mass Loaded Raft

The first step was to determine the new mode shapes. One needs to note that with a non-uniform mass distribution the mode shape is the shape of the displacements or the accelerations of the beam while the modal forces are related to the mass distribution and the modal acceleration thus:

$$F_m \quad = \quad \rho \, \ddot{\phi}_m \qquad\qquad 7.7.1$$

where ρ is the mass distribution and ϕ_m is the displacement modal shape. Orthogonality requires that:

$$\int F_m \, \phi_n \quad = \quad 0 \qquad if \quad m \neq n \qquad\qquad 7.7.2$$

and

$$\int \rho \, \phi_m \, \phi_n \quad = \quad 0 \qquad if \quad m \neq n \qquad\qquad 7.7.3$$

7.7.1 Results and Discussion

The effect of the mass loading of the modes of the raft can be summarised by noting that as the added mass is on the centre-line of the raft one would not expect to see any significant change in the anti-symmetric modes. These include the rigid-body roll mode, the torsion mode and the anti-symmetric U-mode. In addition the rigid body heave mode, surge mode and sway mode should be unchanged while the rigid body pitch mode and yaw mode will rotate about the new centre of gravity. The main change will be to the symmetric U-mode and this could only be determined by a new set of system identification tests as described in section 7.1.

When these system identification tests were repeated for the mass loaded raft a number of unexpected difficulties were encountered. The first of these was that the mode shapes seemed to have changed significantly and the new results were not in line with the predicted changes. This problem was eventually identified as being due to the introduction of the complementary cross-over gap and accelerometer filters. Tests indicated that there were significant errors in the dynamic calibration of the accelerometers and so the gains required to allow for these effects - see Figure 7.5.1. - could not be determined accurately. The problem was overcome, temporarily, by reverting to the gap sensors only for the system identification and the mode shapes were then found to be very close to what had been predicted. The mode shape of the symmetric-U mode had, in fact, changed as indicated in Figure 7.7.1.1 but the other mode shapes remained unchanged or changed by virtue of the shift in the centre of gravity of the raft.

When the correct mode shapes were used in the modal transformation matrix the performance of the mass loaded raft was found to be very similar to that of the bare raft. The response to the simulated sea-way motion is shown in Figure 7.7.1.2.

Again it will be seen that in the absence of the sea-way control there is a large lateral surge displacement of the raft with the raft hitting its end stops. When the sea-way control is included there is again a dramatic decrease in this surge response as shown by the second curve in Figure 7.7.1.2.

The results of the second, more complex, demonstration are shown in Figure 7.7.1.3. Here the raft is directly subject to an external disturbance while its base is being driven by the hydraulic actuator. It will be seen that in the absence of the sea-way control the external disturbance is superimposed on the response of the raft to the movement of its base. When the sea-way control is activated the raft motion disappears but the response to the external disturbance remains unchanged.

Figure 7.7.1.1: Symmetric U-Mode Shape (Mass Loaded)

Fundamental Bending Mode graph with axes Mode Displacement (y-axis, -0.6 to 0.6) and Raft Location (x-axis, 1 to 8). Two curves labelled "Unloaded" and "Mass Loaded".

Figure 7.7.1.2: Suppression of Seaway Motion

Suppression of Sea-Way Motion: Mass Loading graph with axes Surge Displacement (mm) (y-axis, -3.0 to 2.0) and Time (seconds) (x-axis, 0 to 10). Legend: No Seaway Control (solid), Seaway Control (dashed).

Figure 7.7.1.3: Lateral Response to Seaway & Vibration Inputs

In summary these demonstrations indicate that the response of the mass loaded raft to simulated sea-way motions are very similar to the response of the bare raft as seen by comparing these results with Figures 7.3.1 and 7.3.1.5. The raft is stiff to low frequency inputs from the moving base while remaining soft to disturbances generated directly on the raft.

8 Accomplishments

8.1 Significance of Results

We have designed and tested very robust, wideband tri-axial electromagnetic actuators, each fitted with sensors to measure position, acceleration and force in the three orthogonal axes. 16 of these actuators have been made and are used, in two rows of 8 actuators, to support a bare model raft structure. This raft is illustrated in Figure 3.2.1. The sensor information is fed into a digital signal processing system (DSP), which we also designed and built, and the algorithms in the DSP are used to define the drive currents to the 5 electromagnets in each actuator to generate the required control forces. The DSP system consists of a series of identical cards which are coupled together, via a Futurebus+ backplane, to form the completely modular system. Each card has a peak processing capability of 300 MFlops and 12 of these cards can be accommodated in a rack along with a 13th card for communications with other racks or with a host PC. The control algorithms are developed in a very high level language - MATLAB. The translation of the MATLAB code into efficient DSP assembler code is greatly simplified by a virtual machine operating system we have also developed for this DSP system.

Prior to completing and commissioning this model raft, we carried out a series of simulations of the performance we might expect. These simulations assumed that the raft was pin jointed and that the forces would be applied directly at the pin positions - a rather idealised model. These simulations indicated that one should be able to achieve the following results for the bare raft:

(1)   20 dB suppression of the response to sea-way motions;

(2)   the elimination of the mount resonance;

(3)   20 dB suppression of lowest frequency raft resonance;

(4)   40 dB suppression of the acoustic radiation associated with the lowest frequency raft resonance by the use of more complex control force distributions;

(5)   up to 60 dB attenuation of higher frequency (above 40 Hz) vibration transmission.

In the event the actual raft dynamics was considerably more complex than that assumed in the simulations. However, the measured results have, to date, been nearly as good as or, in some cases, actually better namely:

(1)   30 dB suppression of the response to sea-way motions;

(2)   the elimination of the mount resonance;

(3)   20 dB suppression of lowest frequency raft resonance;

(4)   20 - 30 dB suppression of the acoustic radiation associated with the lowest frequency raft resonance by the use of more complex control force distributions.

The more complex force distributions, used to date, are still some way from the optimum distributions, so we expect, ultimately, to pick up an additional 10 dB.

8.2 Implications for Future Design

The implications for future design are:

(1)   the ability to operate a broad band control system on an unprecedented scale which is steadily working up to 264 inputs and 80 outputs;

(2)   the ability to control the raft, via its normal modes, so that one can selectively choose the type of control one wishes to use for each mode;

(3)   the control of modes with complex force distributions which are optimized to meet
        additional but different criteria - low acoustic radiation efficiency in this case;

(4)   the ability to achieve all these goals simultaneously;

(5)   the design of a very robust, wide bandwidth tri-axial electromagnetic actuator with sensors
        to measure position, acceleration and force;

(6)   the design of a modular DSP electronics that can be extended up to 25 GFlops;

(7)   the design of an operating system which greatly simplifies the task of producing and
        testing efficient assembly code for the DSP.

Taken together, these techniques allow the designer to tailor the machinery raft dynamics to meet all
his requirement simultaneously - he need no longer be forced to contend with unsatisfactory
compromises.

9 Future Directions - Phase 3

In the course of the Phase 2 experiments, especially at the last stages, a number of lessons were learnt that need to be carried forward to Phase 3. These are:

(1)   the basic low frequency control strategy works well but the performance is limited by the precision of the input to and output data from the DSP system;

(2)   there is an urgent need for dynamic in-situ calibration techniques to be developed as static calibration cannot be relied on for all the sensors. This will require all 16 sensors from each actuator to be fed into the DSP system rather than 14 sensors per actuator as used in Phase 2. This will require 2 further DSP cards which were not allowed for in the Phase 3 proposal;

(3)   there is a need for more accurate control of the electromagnets both to reduce high frequency noise and to allow the raft mount resonance to be softened;

(4)   there is an absolute need to correct the fault in the daughter cards so that the full power of the DSP system can be utilised;

(5)   there is a need for simpler access to the actuators not only for repair but also to allow a progressive build up in the number of actuators from 16 to 32.

In addition, it would be desirable, if time permits, to instrument and make measurements of the accelerations in the base structure, on which the raft is mounted, to provide independent measurements of the reduction in the force transmitted into the base structure.

10 References

AST/MRC/12/280/TD	DSP (300Mf) Card Defining Specification Issue 2 March 94
AST/MRC/8.0/353/TD	DSP I/O Interface (300Mf) Defining Specification Issue 2 March 93
AST/MRC/8.0/425/TD	Virtual Machine & Library Routines Defining Specification Issue 2 March 94
AST/MRC/14/600/MP	Task Responsibilities Defining Specification Issue 4 February 93
AST/MRC/8.0/611/EN	H.F. Control Algorithm Design Notes Issue 1 January 93
AST/MRC/8.0/615/EN	DSP Design Rules Issue A January 93
AST/ERC/12.3/630/TD	Signal Conditioning & Interconnection Defining Specification Issue 2 March 94
AST/ERC/12.2/631/EN	Gap Sensors Design Notes Issue 2 March 94
AST/ERC/12.0/634/TD	Actuators Defining Specification Issue 2 March 94
AST/ERC/12.4/637/TD	Power Amplifiers Defining Specification Issue 2 March 94
AST/ERC/12.0/638/TD	Raft Defining Specification Issue 2 March 94
AST/ERC/12.2/639/TD	Sensors Defining Specification Issue 2 March 94
AST/ERC/12.2/640/EN	Flux Sensors Design Notes Issue 2 March 94
AST/MRC/14/646/TD	Experimental Programme Defining Specification Issue 1 February 93
AST/MRC/12.5/672/TR	DSP Manufacture Plans Design Notes Issue 2 March 94
AST/MRC/13/673/TR	Virtual Machine Design Notes Issue 1 March 93
AST/MRC/12/674/TR	DSP (300Mf) Card Design Notes Issue 1 March 93
AST/ERC/12.4/684/EN	PWM Power Amplifiers Design Notes Issue 2 March 94
AST/ERC/12.2/685/EN	Flux Sensors Design Notes Issue 2 March 94
AST/ERC/12.1/687/EN	Actuators Design Notes Issue 2 March 94
AST/ERC/12.2/688/EN	Vibration Sensors Design Notes Issue 2 March 94
AST/ERC/11.0/689/TD	Control Theory & Algorithms Defining Specification Issue 2 March 94
AST/ERC/12.4/692/EN	Linear Power Amplifiers Design Notes Issue 2 March 94
AST/MRC/12.0/764/TD	DSP System Defining Specification Issue 1 March 94
AST/MRC/14/771/TD	Experimental Programme Defining Specification Issue 1 July 93
MRC/AST/FAR/11.0	Control Theory Report Sept 93

Distribution

1 - 2	Director, Advanced Research Projects Agency	ARPA, 3701 N Fairfax Dr. Arlington, Va. 22203-1714 DODAAD Code: HR0011
3 - 5	Scientific Officer	ONR, Code 452, 800 North Quincy St, Arlington Va. 22217-5000
6	Administrative Contracting Officer	ONR, 495 Summer St., Room 103, Boston, Ma. 02210-2109
7	Director, Naval Research Laboratory	NRL, Code 2627, Washington, DC. 20375 DODAAD Code: N00173
8 - 9	Defense Technical Information Centre	DTIC, Building 5, Cameron Station, Alexandria, Va. 22304-6145

REPORT DOCUMENTATION PAGE	1. REPORT NO. ARPA/S30M2/94/01	2.	3. Recipient's Accession No.
4. Title and Subtitle Advanced Submarine Technology - Project M Phase 2: Control Experiments and Simulations			S. Report Date March 1994
			6.
7. Author(s) F.A. Johnson			8. Performing Organisation Rept. No. MRC/AST/FSR/21
9. Performing Organization Name and Address GEC- Marconi Research Centre GEC-Marconi Limited Great Baddow, Chelmsford Essex, U.K., CM2 8HN			10. Project/Task/Work Unit No. S30M2 Phase 2
			11. Contract(C) or Grant(G) No. (c) N00014-93-C-0063 (G)
12. Sponsoring Organization Name and Address ARPA 3701 N Fairfax Dr., Arlington Va. 22203-1714 ONR 800N Quincy St, Arlington Va. 22217-5000			13. Type of Report & Period Covered Interim Jan/93-Jan/94
			14.

15. Supplementary Notes

16. Abstract (Limit: 200 words)

The aim of Project M is to demonstrate that the use of broad band active control, with a distributed array of electromagnetic actuators, can not only support a machinery raft but also provide a very high degree of isolation from the hull. In addition the raft can be adequately restrained, to maintain accurate alignment under heavy sea-way or manoeuvring motions, and longer wavelength structural resonances in the raft can be suppressed. This would allow larger lighter rafts can be used to accommodate machinery items without an increase in the acoustic radiation hazard.

This document brings together a description of the work on the control theory, the simulations, the hardware upgrades and the results of the experimental programme which validated the theoretical predictions.

17. Document Analysis a. Descriptors
Active control, Active suspension, Active vibration isolation, Actuators, Digital signal processors, Flexible support, Machinery isolation, Magnetic levitation, Magnetic suspension, Submarines, Vibration control.

b. Identifiers/Open-Ended Terms

c. COSATI Field/Group

18. Availability Statement Unlimited	19. Security Class (This Report) Unclassified	21. No. Of Pages 94
	20. Security Class (This Page) Unclassified	22 Price ....

(see ANSI-Z39.18)   OPTIONAL FORM 272 (4-77)